

Validation of Turbulent Transport Models on Alcator C-Mod and ASDEX Upgrade

by

Alexander James Creely

B.S.E., Princeton University (2014)

Submitted to the Department of Nuclear Science and Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Applied Plasma Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2019

© Massachusetts Institute of Technology 2019. All rights reserved.

Author

Department of Nuclear Science and Engineering

February 8, 2019

Certified by

Anne E. White

Cecil and Ida Green Associate Professor in Nuclear Engineering

Thesis Supervisor

Certified by

Garrard Conway

Senior Staff Scientist, Max Planck Institute for Plasma Physics

Thesis Reader

Accepted by

Ju Li

Battelle Energy Alliance Professor of Nuclear Science and Engineering

Chairman, Department Committee on Graduate Theses

Validation of Turbulent Transport Models on Alcator C-Mod and ASDEX Upgrade

by

Alexander James Creely

Submitted to the Department of Nuclear Science and Engineering
on February 8, 2019, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Applied Plasma Physics

Abstract

This thesis developed hardware and analysis techniques to measure two validation constraints experimentally, and then applied these constraints in the validation of plasma turbulent transport models on two tokamaks, Alcator C-Mod and ASDEX Upgrade, resulting in both greater physics understanding of multi-scale turbulent interactions and greater confidence in predictions for future fusion devices.

On the path toward the clean, sustainable, and safe energy of a fusion power plant, experiment and modeling each contribute something unique. Before one can in good faith use plasma turbulent transport models to explain turbulent dynamics or predict machine performance, however, one must ensure that these models can correctly reproduce experimentally measured conditions on existing devices. Validation, the process of determining how accurately a model represents reality, has thus become a key endeavor in fusion energy research.

First, this thesis developed an analysis technique to measure the electron perturbative thermal diffusivity based on tracking the propagation of heat pulses generated by partial sawtooth crashes. In addition, correlation electron cyclotron emission (CECE) hardware was constructed on both Alcator C-Mod and ASDEX Upgrade, and analysis techniques were derived, in order to measure turbulent electron temperature fluctuations.

These validation constraints were applied to two turbulent transport models, the nonlinear gyrokinetic model and the quasi-linear gyrofluid model. In particular, these constraints were used to study the importance of multi-scale turbulent effects (due to coupling between ion- and electron-scales) in correctly modeling plasma behavior. The gyrokinetic codes GYRO and GENE were validated on Alcator C-Mod and ASDEX Upgrade respectively, using both constraints developed in this thesis as well as ion and electron heat fluxes from power balance, revealing that in some cases ion-scale simulations are sufficient to match experimental constraints, while in other cases multi-scale effects are important.

To investigate this discrepancy, a novel type of validation study was performed

with the gyrofluid code TGLF, including many discharges from both machines. This study resulted in two physical criteria that determine when multi-scale effects are important, and when ion-scale simulations are sufficient to model the plasma behavior, shedding light on the physical phenomena that govern the importance of multi-scale turbulent effects.

Thesis Supervisor: Anne E. White

Title: Cecil and Ida Green Associate Professor in Nuclear Engineering

Acknowledgments

While the work of a doctoral thesis is to some extent a personal endeavor, it is by no means accomplished alone. I would never have made it to this point without the generous support of my colleagues and friends throughout the process, and many thanks are in order.

First and foremost, I would like to thank my advisor, Professor Anne White, for her incredible support throughout my graduate student career. I have learned so much from her both scientifically and professionally, and I don't think that there is anything else that someone could ask for in an advisor.

I would also like to thank Dr. Garrard Conway and Dr. Simon Freethy for their mentorship and friendship during my time at ASDEX Upgrade, and for our weekly meetings. They were both invaluable resources of scientific knowledge and of advocacy at ASDEX Upgrade, and good friends during my many visits. I would also like to thank Dr. Garrard Conway for serving as my thesis reader.

I would like to thank Dr. Martin Greenwald for serving on my thesis committee, as well as for many helpful suggestions along the way.

Dr. Nathan Howard taught me an incredible amount about transport modeling and about how to run various codes on Alcator C-Mod. I would also like to thank him for running GYRO for some of my validation studies.

I would like to thank Dr. Tobias Görler for teaching me to run GENE on ASDEX Upgrade, and also for valuable discussions about simulations and their results.

Dr. Choongki Sung was my predecessor on the CECE project at MIT, and helped guide me through the first year of my PhD. His thesis also taught me a lot about CECE diagnostics.

More generally, I would like to thank everyone who helped gather the data that I used in this thesis. On Alcator C-Mod, data from Dr. Amanda Hubbard, Dr. John Rice, and Norman Cao was particularly useful. I would also like to thank the Alcator C-Mod Team in general and the PSFC at large for running Alcator C-Mod and allowing me to work with such an incredible machine.

Similarly, on ASDEX Upgrade, I would like to thank Dr. Giovanni Tardini, Dr. Tim Happel, Dr. Pascale Hennequin, Dr. Rachael McDermott, Dr. Matthias Wilensdorfer, and Dr. Tomas Odstrcil for their help with data and hardware, as well as the ASDEX Upgrade Team for running an amazing facility.

On the modeling front, I would like to thank Dr. Gary Staebler for his help with running TGLF, Pablo Rodriguez-Fernandez for working with me on VITALS (among many other things), and Dr. Orso Meneghini for helping me to use OMFIT.

In terms of hardware, I would like to thank David Terry, Willy Burke, Rick Lecacorvi, Bill Parkin, and Dr. Jim Irby for working with me during the design and fabrication of the ASDEX Upgrade CECE system. Without their design and engineering expertise, the final diagnostic would surely not have been the amazing system that it is.

In addition, I would like to thank Professor Katsumi Ida for advising me during my time at the Large Helical Device and Dr. Chuck Kessel for teaching me to run TSC.

My time at the PSFC surely would not have run so smoothly or have been nearly as much fun without Valerie Censabella, Anastasia Alexandris, and Dr. Ilavenil Subbiah, with whom I had a wonderful time working.

Finally, I would like to thank all of my friends for their wonderful support throughout my time in graduate school. In particular, I would like to thank Shen Ning for her never-ending laughs and love.

In addition to support from various individuals and teams, this thesis has benefitted from funding, logistic, and computational support from many sources.

The experimental work on Alcator C-Mod was funded by the US Department of Energy under grant DE-FC02-99ER54512-CMOD. In addition the US DOE funded the development of the CECE diagnostics for both Alcator C-Mod and ASDEX Upgrade under grants DE-SC0006419 and DE-SC0017381.

The experimental work on ASDEX Upgrade was carried out within the framework of the EUROfusion Consortium and received funding from the Euratom research and training program for 2014-2018 under grant agreement No 633053. Note that the

views and opinions in this work do not necessarily reflect those of the European Commission.

GYRO and GENE simulations were performed with MP224 Repo at the National Energy Research Scientific Computing Center (NERSC), a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

TGLF simulations used resources of the MIT PSFC parallel AMD Opteron/Infiniband Loki cluster and the MIT PSFC Engaging Cluster at the Massachusetts Green High Performance Computing Center (MGHPCC).

My personal stipend and tuition funding came from three sources. In my first year, I was funded by the MIT Charles A. Piper (1935) Presidential Fellowship. In my second year, I was funded by the US DOE under one of the CECE grants, DE-SC0006419. In my third through fifth years, I was funded by the US Department of Defense and the Air Force Office of Scientific Research under the National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a.

During my time in Japan at the National Institute for Fusion Science, I was funded by the MIT International Science and Technology Internship (MISTI) Program. The Large Helical Device itself is partly supported by a Grant-in-Aid for Scientific Research of the Japan Society for the Promotion of Science (JSPS) (No. 15H02336). The work on LHD was also partly supported by the National Institute for Fusion Science grant administrative budget (No. NIFS10ULHH021).

Finally, I received funding for various travels and professional development from the Hugh Hampton Young Memorial Fund Fellowship at MIT.

Contents

Glossary of Terms	12
List of Symbols	18
List of Important Equations	23
List of Publications	25
1 Introduction	29
2 Background and General Concepts	37
2.1 Turbulence and Transport in Tokamaks	37
2.2 Turbulent Transport Models	45
2.3 Tokamaks and Standard Diagnostics in this Thesis	56
2.3.1 Alcator C-Mod	57
2.3.2 ASDEX Upgrade	58
2.4 Validation and Validation Constraints	59
2.4.1 Power Balance Heat Fluxes	62
2.4.2 Electron Perturbative Thermal Diffusivity	69
2.4.3 Electron Temperature Fluctuations	74
3 Perturbative Thermal Diffusivity From Partial Sawtooth Heat Pulses	81
3.1 Partial Sawteeth	82
3.2 Perturbative Diffusivity Calculation Method	85
3.3 Comparison to Full Sawtooth Heat Pulses and Modulated Electron Cyclotron Heating	94
3.4 Experimental Observations	98
3.4.1 Alcator C-Mod	99
3.4.2 Cross-Machine Correlations on ASDEX Upgrade and Alcator C-Mod	109
4 Electron Temperature Fluctuations from Correlation Electron Cy- clotron Emission	115
4.1 CECE Radiometers	116
4.1.1 Alcator C-Mod Hardware	117
4.1.2 ASDEX Upgrade Hardware	120
4.2 Data Analysis Techniques	125
4.2.1 Calculation of Temperature Fluctuation Levels	126
4.2.2 Testing of Data Analysis Techniques	131

4.3	Experimental CECE Measurements	139
4.3.1	Alcator C-Mod	140
4.3.2	ASDEX Upgrade	147
5	Validation of Gyrokinetic Codes	153
5.1	Validation of GYRO on Alcator C-Mod	154
5.1.1	L- and I-mode Validation	155
5.1.2	Validation of Ion- and Multi-Scale GYRO Simulations with Per- turbative Diffusivity	175
5.2	Validation of GENE on ASDEX Upgrade	179
5.2.1	Experimental Conditions of ASDEX Upgrade Discharge 33585	180
5.2.2	Inner Radius Validation Study Perturbative Thermal Diffusiv- ity Measurements	182
5.2.3	Outer Radius Validation Study with CECE Measurements . .	202
5.3	General Validation Observations and Conclusions	210
6	Validation of TGLF in Multi-Machine Study	215
6.1	TGLF Applied to L- and I-mode Validation Study	216
6.2	Methodology and Plasmas for Multi-Machine Study	221
6.3	Results of Multi-Machine Study	228
6.3.1	Validation Results	228
6.3.2	Validation Metric and Criteria for the Importance of Multi- Scale Effects	232
6.3.3	Profile Prediction	241
7	Conclusions	247
A	History of Perturbative Thermal Diffusivity Measurements	257
B	Calculation of Numerical Prefactor for Perturbative Thermal Diffu- sivity	261
C	Derivation of CECE Equations	267
C.1	Frequency Domain Derivation of Total Temperature Fluctuation . . .	267
C.2	Time Domain Derivation of Total Temperature Fluctuation	272
C.3	Derivation of Uncertainty	275
C.4	Derivation of Sensitivity Limit	276
C.5	Optical Depth Derivation	278
D	The Ricci Validation Metric	281
E	Density Fluctuation Measurements with Reflectometry on the Large Helical Device (LHD)	285
E.1	Introduction to Density Fluctuation Measurements	286
E.2	Experimental Set-Up on LHD	288
E.3	Reflectometry Analysis Methods	290

E.4	Density Fluctuation Profile Measurements	297
F	MQ1 Tokamak Design using the Tokamak Simulation Code (TSC)	305
F.1	The MQ1 Tokamak	307
F.2	Machine and Coil Geometry	309
F.3	Discharge Evolution	312
F.4	Predicted Peak Plasma Performance	314
G	InterPol Interferometer Polarimeter	317
G.1	Interferometry and Polarimetry	318
G.2	Proposed InterPol Diagnostic	321
G.3	Predicted Signals and Measurement Capabilities	325
H	Specifics of Data Analysis and Simulation Workflows	333
H.1	Perturbative Diffusivity Analysis	333
H.2	CECE Data Analysis	336
H.3	TRANSP	339
H.4	GENE	341
H.5	TGLF and VITALS	343
	Bibliography	351

Glossary of Terms

Tokamak Physics, Operation, and Control

MHD	Magneto-Hydro-Dynamics - A highly simplified model of magnetized plasmas used for macroscopic stability calculations.
Grad-Shafranov Equation	An equation describing the MHD equilibrium of toroidal systems.
Shafranov Shift	A shift of the magnetic axis from the geometric plasma center due to finite plasma pressure.
TAE	Toroidal Alfvén Eigenmode - A coherent plasma oscillation associated with the strength of the magnetic field and the toroidicity of the system.
L-mode	Low confinement mode. The standard mode of operation in an auxiliary heated tokamak.
H-mode	High confinement mode. A mode of tokamak operation where, given sufficient heating power, the plasma edge self organizes into a large pedestal in both temperature and density.
I-mode	Improved confinement mode. A mode of tokamak operation where, in a range of heating power, the plasma edge self organizes into a large pedestal in temperature but not in density.
Core	The center of the tokamak plasma.
Pedestal/Edge	The outer region of the plasma, where steep gradients are often found, and some core turbulence models break down.
Separatrix	The edge of the plasma in a diverted machine.
SOL	Scrape Off Layer - The region outside of the last closed flux surface, where magnetic field lines connect to a solid surface.
X-point	The magnetic null in a diverted plasma.
ELM	Edge Localized Mode - A periodic burst of energy and particles released from the pedestal during standard H-mode operation.
RMP	Resonant Magnetic Perturbation - An oscillating magnetic field applied to the edge of a plasma in order to increase transport and avoid ELMs.
Impurity	A plasma species other than the main ion species.
TF	Toroidal Field (Coil) - The magnetic coils that generate the dominant toroidal field in the tokamak.
OH	Ohmic Heating (Coil) - Also called the Central Solenoid. The magnetic coil in the center of the tokamak used for current drive, ohmic heating, and shaping.

EF	Equilibrium Field (Coil) - The magnetic coils that generate the poloidal field necessary for radial stability and shaping of the plasma.
PID	Proportional-Integral-Derivative - A type of controller that actuates based on components proportional to the signal, its derivative, and its integral.
Tokamak Diagnostics and Heating	
ECE	Electron Cyclotron Emission - Radiation emitted from the plasma due to the gyration of electrons around the magnetic field. Used to measure electron temperature.
GPC	Grating Polychrometer - One type of diagnostic used to detect ECE based on a diffraction grating.
Radiometer	Another type of diagnostic used to detect ECE, based on a series of bandpass filters.
Interferometry	A laser based diagnostic that measures line-integrated density.
FIR	Far Infrared (Interferometry) - A particular wavelength of laser used for interferometry.
Polarimetry	A laser based diagnostic that measures the line-integrated magnetic field parallel to the beam propagation path.
PCI	Phase Contrast Imaging - A laser based diagnostic that measures line-integrated density fluctuations in the plasma.
TS	Thomson Scattering - A laser scattering diagnostic used to measure both electron temperature and density.
CXRS	Charge Exchange Recombination Spectroscopy - Also known as CXS (Charge Exchange Spectroscopy). A diagnostic used to measure ion temperature and plasma rotation. Requires neutral beam injection.
XICS	X-Ray Imaging Crystal Spectrometer - A diagnostic used to measure ion temperature and plasma rotation. Requires injection of a trace impurity, typically Argon or Krypton.
Doppler Reflectometry	A microwave diagnostic used to measure density and density fluctuations based on the reflection of an injected wave. Sometimes also called DBS (Doppler Backscattering).
BES	Beam Emission Spectroscopy - A diagnostic used to measure density fluctuations. Requires neutral beam injection.
NBI	Neutral Beam Injection - Also known as NBH (Neutral Beam Heating). A beam of high energy neutral particles injected into the tokamak for heating and fueling of the plasma.

ECH	Electron Cyclotron Heating - Also known as ECRH (Electron Cyclotron Resonance Heating). The launching of waves at the electron cyclotron resonance frequency for plasma heating.
ICRH	Ion Cyclotron Resonance Heating - Also known as ICRF (Ion Cyclotron Range of Frequencies). The launching of waves at the ion cyclotron resonance frequency (or its harmonics) for plasma heating. Often utilizes a minority species for better absorption.
RF	Radio Frequency - A term sometimes used to collectively refer to ECH and ICRH.
Line-Averaged	A measurement that integrates a plasma property along a line of sight or a beam path.
Perturbative Diffusivity Measurement	
Heat Pulse	A transient increase in temperature that propagates through the plasma. Can be generated in a variety of ways, including sawtooth crashes and modulated ECH.
Sawtooth Crash	An MHD event in the plasma core involving magnetic reconnection at the $q=1$ surface. Causes macroscopic transport of heat and particles out of the plasma core.
Partial Sawtooth Crash	Similar to a full sawtooth crash, but only involving an annular region of the plasma core. Leads to significantly smaller heat redistribution.
Inversion Radius	The radius at which the temperature and density profiles do not change during a sawtooth crash.
Mixing Radius	The radius out to which the sawtooth crash flattens the temperature and density profiles.
Ballistic Effect	Non-diffusive heat transport beyond the mixing radius associated with the sawtooth crash.
CECE Diagnostic and Analysis	
CECE	Correlation Electron Cyclotron Emission - A specialized ECE radiometer that correlates closely spaced measurement volumes in order to detect small amplitude temperature fluctuations.
O-mode	Ordinary Mode - Light propagating with its electric field parallel to the background magnetic field.
X-mode	Extraordinary Mode - Light propagating with its electric field perpendicular to the background magnetic field.
RF Section	Radio Frequency Section - The high frequency electronics section of the CECE diagnostic, operating in the 100s of GHz.
IF Section	Intermediate Frequency Section - The lower frequency electronics section of the CECE diagnostic, operating in the single to 10s of GHz.

YIG	Yttrium Iron Garnet - A type of tunable bandpass filter.
Sensitivity Limit	The level below which a signal is indistinguishable from noise.
Synthetic Data	Artificially generated fluctuation data with added noise used for testing of analysis routines.
Doppler Broadening	An increase in the ECE emission volume width due to a Doppler shift in the emitted ECE frequency.
Relativistic Broadening	An increase in the ECE emission volume width due to the relativistic mass increase and its effect on the cyclotron frequency.
Turbulence	
Drift Wave	The dominant cause of turbulence in tokamak plasmas. Generated through an interplay between pressure gradients and diamagnetic currents in the plasma.
ITG	Ion Temperature Gradient (Mode) - A turbulent mode driven unstable by gradients in the ion temperature.
TEM	Trapped Electron Mode - A turbulent mode involving trapped electrons and driven unstable by electron temperature and density gradients.
ETG	Electron Temperature Gradient (Mode) - A turbulent mode driven unstable by gradients in the electron temperature.
Zonal Flow	Ordered flows that develop in a tokamak perpendicular to the pressure gradients. Zonal flows interact in complex ways with various turbulent modes.
Flux-Gradient Relationship	The dependence of the turbulent heat flux on the driving gradient.
Linear Growth Rate	The characteristic growth rate of a single unstable turbulent mode, ignoring interactions with other modes.
Dominant Mode	The turbulent mode with the largest linear growth rate at a given wavenumber.
Sub-Dominant Mode	A turbulent mode that is unstable at a given wavenumber, but has a smaller growth rate than another mode.
Turbulence Models	
Gyrokinetic Model	A reduced kinetic model of plasma turbulence based on the small size of the gyro-radius compared to the gradient scale lengths of the machine, as well as other assumptions.
Gyrofluid Model	A fluid model based on taking the fluid moments of the gyrokinetic equations.
Coppi-Tang Model	A highly simplified analytic model of turbulent heat and particle transport in tokamak plasmas.

Linear	A model that only calculates the linear growth rates of turbulent modes.
Nonlinear	A model that calculates the nonlinear saturated state of turbulence.
Quasi-Linear	A model that calculates linear growth rates, and then uses these to approximate the saturated state of the turbulence based on some saturation rule.
Saturation Rule	A rule used to translate linear growth rates into a nonlinear saturated state, typically generated using the results of previous nonlinear simulations.
GYRO	A nonlinear gyrokinetic simulation code.
GENE	Gyrokinetic Electromagnetic Numerical Experiment - Another nonlinear gyrokinetic simulation code.
TGLF	Trapped Gyro-Landau Fluid - A quasi-linear gyrofluid simulation code.
Electrostatic	A simulation that allows only the plasma potential, and not the magnetic field, to fluctuate.
Electromagnetic	A simulation that allows both the plasma potential and the magnetic field to fluctuate.
Ion-Scale	A simulation that resolves down to the ion gyro-radius.
Electron-Scale	A simulation that resolves only around the electron gyro-radius.
Multi-Scale	A simulation that resolves all the way from larger than the ion gyro-radius down to the electron gyro-radius.
Local	A simulation of a single radius in the tokamak, with a domain of a single flux-tube.
Global	A simulation of an extended radial region of the tokamak.
	Validation
Validation	“The process of determining the degree to which a model is an accurate representation of the real world from the perspective of the intended uses of the model.” [1]
Validation Constraint	Also known as a quantity of interest. An experimentally measured quantity that one can compare to the outputs of a simulation in order to validate the simulation.
Validation Metric	A means of quantifying the overall agreement between experimental measurements and the results of a simulation.
Validation Study	A reasonably self-contained effort to validate one or a set of codes on one or a set of plasmas.
Model Selection	The task of selecting the appropriate model for a particular task.

Fortuitous Agreement	A result where a model will give the correct answer but for the wrong physical reason. Additional constraints reduce the likelihood of fortuitous agreement.
Verification	The process of checking that a particular code has accurately implemented a given model.
Power Balance	The calculation balancing heating and loss terms in order to infer the turbulent heat flux in a tokamak plasma.
Synthetic Diagnostic	A method of analyzing the output of simulations in order to compare to real experimental measurements, taking into account the limitations of the real diagnostic.
Heat-Flux Matched Simulation	A simulation that best matches the inferred experimental heat fluxes, which is then used to compare to other validation constraints.
Other Codes	
EFIT	A code used to reconstruct the magnetic equilibrium of a tokamak plasma. Used widely at several machines around the world.
CLISTE	Another code used to reconstruct the magnetic equilibrium of a tokamak plasma. Used as ASDEX Upgrade.
TRANSP	A power balance solver used widely around the world.
MDSplus	The data storage system at Alcator C-Mod.
TGYRO	A wrapper for TGLF and GYRO, used to predict temperature and density profiles.
NEO	A code used to calculate neoclassical transport levels.
VITALS	Validation via Iterative Training of Active Learning Surrogates - A wrapper for TGLF, used to vary many input parameters within uncertainty in order to best match a set of validation constraints.
OMFIT	A user interface allowing access to many codes and analysis routines.
EPED	A code used to predict the pedestal of H-mode plasmas.
LHDGauss	A ray tracing code used to track the propagation of waves in LHD.
TSC	Tokamak Simulation Code - A code used to design the poloidal field coil set for tokamaks. Solves for MHD stability and uses the Coppi-Tang model to approximate the plasma behavior.

List of Symbols

General Physics Parameters

c	Speed of light
ϵ_0	Permittivity of free space
μ_0	Permeability of free space
e	Fundamental charge
t	Time
t'	Retarded time
Ω	Solid angle
ω	Angular frequency
ν	Frequency
\vec{v}	Velocity
\vec{a}	Acceleration
β_r	Relativistic correction beta ($\beta_r = v/c$)
\vec{B}	Magnetic field
\vec{E}	Electric field
∇	Gradient in physical space
∇_v	Gradient in velocity space
V	Volume
m_s	Mass of species s

Mathematical Notation

\tilde{x}	Fluctuations in quantity x
\bar{x}	Line averaged value of quantity x
x	Scalar x
\vec{x}	Vector \vec{x}
$\overline{\overline{X}}$	Tensor $\overline{\overline{X}}$

Kinetic Plasma Description

f_s	Distribution function of species s
n_s	Number density of species s
T_s	Temperature of species s
Z_s	Charge number of species s
q_s	Charge of species s ($q_s = Z_s e$)
p_s	Pressure of species s ($p_s = n_s T_s$)
\vec{J}_s	Current carried in species s
$C(f_s, f_{s'})$	Collision operator in Boltzmann Equation

Fluid Plasma Description

S_j	Volumetric power density in quantity j (MW/m ³)
Q_j	Power flux in quantity j (MW/m ²)
P_j	Total power in quantity j (MW)
C_s	Volumetric energy exchange due to collisions of species s
$\overline{\overline{P}}_s$	Pressure tensor of species s
\vec{u}_s	Fluid velocity of species s
\vec{q}_s	Conductive heat flux vector of species s
W_s	Stored energy density of species s ($W_s = (3/2)n_s T_s$)
U_j	Total stored energy of species s ($U_s = \int W_s dV$)
$S_{conv,s}$	Local heating due to convection of species s ($S_{conv,s} = \vec{u}_s \cdot \nabla W_s$)
$S_{comp,s}$	Local heating due to compression of species s ($S_{comp,s} = -p_s \nabla \cdot \vec{u}_s$)
$S_{cond,s}$	Local heating due to conduction through species s ($S_{cond,s} = \nabla \cdot \vec{q}_s$)
S_{ei}	Local heating due to collisional heat transfer from electrons to ions
S_{RF}	Local heating from radio frequency waves
S_{NBI}	Local heating from neutral beam injection
S_{OH}	Local heating from ohmic dissipation
S_{rad}	Local power loss due to radiation
S_{ion}	Local power loss due to the ionization of neutrals
S_{cx}	Local power loss due to charge exchange reactions
$S_{i,s}$	Local net power gain from the ion particle source
Q_s	Total turbulent and neoclassical heat flux of species s ($Q_s = Q_{cond,s} + Q_{conv,s}$)
χ_s^{PB}	Power balance diffusivity of species s ($\chi_s^{PB} = -\frac{Q_s}{n_s \nabla T_s}$)
χ_s^{pert}	Perturbative diffusivity of species s ($\chi_s^{pert} = -\frac{1}{n_s} \frac{\partial Q_s}{\partial \nabla T_s}$)

General Plasma Parameters

p	Total plasma pressure ($p = \sum_s p_s$)
β	Plasma beta ($\beta = p/(B^2/2\mu_0)$)
v_{T_s}	Thermal velocity of species s ($v_{T_s} = \sqrt{2T_s/m_s}$)
c_s	Ion sound speed ($c_s = \sqrt{2T_e/m_i}$)
ρ_i	Ion Larmor radius, also known as the ion gyro-radius ($\rho_i = v_{T_i}/\omega_{ci}$)
ρ_e	Electron Larmor radius, also known as the electron gyro-radius ($\rho_e = v_{T_e}/\omega_{ce}$)
ω_{cs}	Cyclotron frequency of species s ($\omega_{cs} = q_s B/m_s$)
\vec{v}_E	$E \times B$ velocity of the plasma ($\vec{v}_E = (\vec{E} \times \vec{B})/B^2$)
Z_{eff}	Effective charge of the plasma ($Z_{eff} = (\sum_s Z_s^2 n_s)/n_e$)
ν_{ei}	Electron-ion collision frequency

ν^*	Collisionality (electron-ion frequency normalized by the electron bounce frequency)
ν_{GENE}	Species independent collision frequency used in the GENE code (internally ν_c)
v_A	Alfvén speed ($v_A = (B^2/\mu_0 n_i m_i)^{1/2}$)
f_A	Alfvén frequency ($f_A \approx v_A/(4\pi qR)$)
Tokamak Parameters	
B_t	Toroidal magnetic field (typically quoted on the magnetic axis)
I_p	Total plasma current
P_{aux}	Total auxiliary heating (everything except ohmic heating)
P_{fusion}	Total fusion power
Q	Fusion gain ($Q = P_{fusion}/P_{heat}$)
r	Local minor radius in a tokamak
a	Machine minor radius
R	Machine major radius
q	Safety factor
q_{95}	Safety factor at the flux surface containing 95% of the toroidal flux
κ	Elongation
δ	Triangularity
ρ_{tor}	Square root of the toroidal magnetic flux normalized by the toroidal flux at the last closed flux surface
ρ_{pol}	Square root of the poloidal magnetic flux normalized by the poloidal flux at the last closed flux surface
τ_E	Energy confinement time
$H_{98,y2}$	Confinement factor from 1998 Type II ELMy H-mode ITER scaling
Stellarator Parameters	
r_{eff}	Effective minor radius
a_{99}	Effective minor radius within which 99% of the total plasma stored energy is confined
Turbulence Parameters	
ρ_s	Ion Larmor radius evaluated at the electron temperature ($\rho_s = c_s/\omega_{ci}$)
ρ^*	Normalized ion Larmor radius evaluated at the electron temperature ($\rho^* = \rho_s/a$)
a/L_y	Normalized gradient scale length of parameter y ($a/L_y = (a/y)(\partial y/\partial r)$)
k_{\parallel}	Wavenumber parallel to the magnetic field
k_{\perp}	Wavenumber perpendicular to the magnetic field
k_{θ}	Poloidal wavenumber

k_y	Binormal wavenumber
λ	Fluctuation wavelength
\hat{s}	Safety factor shear ($\hat{s} = r d \ln(q)/dr$)
s_κ	Elongation shear ($s_\kappa = r d \ln(\kappa)/dr$)
s_δ	Triangularity shear ($s_\delta = r d \ln(\delta)/dr$)
γ	Linear growth rate of a turbulent mode at a given wavenumber
ω	Real frequency of a turbulent mode at a given wavenumber
$\gamma_{E \times B}$	$E \times B$ shearing rate
γ_p	Rotation shearing rate
$Q_{GB,s}$	Gyro-Bohm heat flux of species s ($Q_{GB,s} = (T_s^{3/2} n_s / m_s^{1/2})(\rho^*)^2$)
L_r	Radial correlation length of a turbulent mode
α_{nT}	Phase angle between density and temperature fluctuations

Optical Parameters

$I(\nu)$	Intensity of light at frequency ν
τ	Optical thickness, also called optical depth
w_0	Beam width ($1/e^2$ power diameter)
λ	Wavelength
$\Delta\phi$	Interferometer phase shift
α	Faraday rotation angle

Perturbative Diffusivity Calculation

v_{HP}	Velocity of the peak of the heat pulse
α	Radial damping parameter for heat pulse
s	Shafranov shift
A	Pulse amplitude
a_c	Plasma minor radius corrected for elongation ($a_c = \sqrt{\kappa}a$)
τ_d	Heat pulse damping time constant

Statistical Parameters

σ_y	Standard deviation of parameter y
R^2	The amount of variation in one parameter that can be explained by the variation in another parameter
γ_c	Complex coherence function ($\gamma_c = G_{xy} / \sqrt{G_{xx}G_{yy}}$)
γ_{bg}	Background coherence
γ_{xy}	Magnitude of the complex coherence function $\gamma_{xy} = \gamma_c $
G_{xy}	One-sided cross-spectral density function ($G_{ij} = 2F_i^* F_j$)
G_{xx}	Auto-spectral density function
R_{xy}	Cross-correlation function
R_{xx}	Auto-correlation function

F_i	Fourier transform of channel i
n_d	Number of independent ensemble averaging windows
N	Total number of data points
δf	Frequency resolution fo the Fourier transform
b	Bias error
$Re\{x\}$	Real part of x

CECE Electronics and Signal

B_{IF}	Intermediate frequency filter bandwidth
B_{sig}	Signal bandwidth
B_{vid}	Video bandwidth
B_{samp}	Sampling frequency
\tilde{T}_e	Time domain temperature fluctuations
\tilde{N}_j	Time domain radiometer thermal noise in channel j
\tilde{N}_j	Frequency domain radiometer thermal noise in channel j

Validation Parameters

χ_{Ricci}	Ricci validation metric
$\Delta\chi_{Ricci}$	Difference between Ricci metric from ion-scale and multi-scale simulations
R_j	Level of agreement in parameter j for Ricci metric calculation
H_j	Validation hierarchy level of parameter j for Ricci metric calculation
S_j	Uncertainty quantification of parameter j for Ricci metric calculation
$\Gamma_{high/low}$	Criteria for the importance of multi-scale effects based on ratio of high and low wavenumber linear growth rates
$\Sigma_{tem/itg}$	Criteria for the importance of multi-scale effects based on ratio of TEM and ITG linear growth rates

List of Important Equations

Basic Plasma Equations

$$\omega_{cs} = \frac{q_s B}{m_s} \quad (1) \quad Z_{eff} = \frac{\sum_s Z_s^2 n_s}{n_e} \quad (9)$$

$$\omega_{pe} = \sqrt{n_e e^2 / m_e \epsilon_0} \quad (2) \quad \vec{v}_E = (\vec{E} \times \vec{B}) / B^2 \quad (10)$$

$$v_{T_s} = \sqrt{\frac{2T_s}{m_s}} \quad (3) \quad f_A \approx \frac{v_A}{4\pi q R} \approx \frac{B_0}{4\pi q R \sqrt{\mu_0 \sum n_s m_s}} \quad (11)$$

$$c_s = \sqrt{\frac{2T_e}{m_i}} \quad (4) \quad L_y = y(\partial y / \partial r)^{-1} \quad (12)$$

$$\rho_s = \frac{c_s}{\omega_{ci}} \quad (5) \quad \nu^* \approx 0.01 \left(\frac{R_0}{r}\right)^{3/2} \left(\frac{q R_0 n_{20}}{T_k^2}\right) \quad (13)$$

$$\rho_i = \frac{v_{T_i}}{\omega_{ci}} \quad (6)$$

$$\rho_e = \frac{v_{T_e}}{\omega_{ce}} \quad (7) \quad \nu_{GENE} = 2.3031 \cdot 10^{-5} \frac{a \cdot n_{19} \cdot \ln \Lambda}{T_k^2} \quad (14)$$

$$\beta = \frac{p}{B^2 / 2\mu_0} \quad (8) \quad \frac{P_f}{V} \propto \langle p \rangle^2 \propto \beta^2 B_0^4 \quad (15)$$

Perturbative Diffusivity

$$\chi_e^{PB} = -\frac{Q_e}{n_e \nabla T_e} \quad (16)$$

$$\chi_e^{pert} = -\frac{1}{n_e} \frac{\partial Q_e}{\partial \nabla T_e} \quad (17)$$

$$\chi_e^{pert} = 4.2 a_c \frac{v_{HP}}{\alpha} \quad (18)$$

$$v_{HP} = \sqrt{\kappa} \frac{a}{a-s} \left(\frac{dt_{peak}}{dr}\right)^{-1} \quad (19)$$

$$\alpha = 10(a-s) \frac{d}{dr} \log(A) \quad (20)$$

CECE

$$\frac{\tilde{T}}{T} = \sqrt{\frac{2}{B_{IF}} \int_{f_1}^{f_2} \frac{Re\{\gamma_c(f) - \gamma_{bg}\}}{1 - Re\{\gamma_c(f) - \gamma_{bg}\}} df} \quad (21)$$

$$\gamma_c(f) = \frac{G_{xy}(f)}{\sqrt{G_{xx}(f)G_{yy}(f)}} \quad (22)$$

$$\sigma_{\tilde{T}/T} \approx \frac{1}{\tilde{T}/T} \frac{1}{B_{IF}} \sqrt{\sum_i (\sigma_{\gamma_i} \delta f)^2} \quad (24)$$

$$\sigma_{\gamma_c}(f) = \sqrt{\frac{1}{2n_d} (1 - |\gamma_c(f)|^2)^2} \quad (23)$$

$$\tau \sim \int n_e T_e^{m-1} ds \quad (25)$$

Validation

$$\chi_{Ricci} = \frac{\sum_j R_j H_j S_j}{\sum_j H_j S_j} \quad (26)$$

$$\Delta\chi_{Ricci} = \chi_{ion} - \chi_{multi} \quad (30)$$

$$R_j = \frac{\tanh[(d_j - d_0)/\lambda] + 1}{2} \quad (27)$$

$$\Gamma_{high/low} = (\gamma_{high-k}/k_{high}) / (\gamma_{low-k}/k_{low}) \quad (31)$$

$$d_j = \sqrt{\frac{1}{N_j} \sum_{i=1}^{N_j} \frac{(x_{j,i} - y_{j,i})^2}{\Delta x_{j,i}^2 + \Delta y_{j,i}^2}} \quad (28)$$

$$\Sigma_{tem/itg} = \exp\left(\frac{-(\frac{\gamma_{TEM}/k_{TEM}}{\gamma_{ITG}/k_{ITG}} - 1.0)^2}{0.09}\right) \quad (32)$$

$$S_j = \exp\left(-\frac{\sum_i \Delta x_{j,i} + \sum_i \Delta y_{j,i}}{\sum_i |x_{j,i}| + \sum_i |y_{j,i}|}\right) \quad (29)$$

Interferometry and Polarimetry

$$\Delta\phi = \frac{-e^2\lambda}{4\pi c^2 m_e \epsilon_0} \int n_e \left[1 - \frac{3}{2} \frac{T_e}{m_e c^2}\right] dl \quad (33)$$

$$\alpha = \frac{e^3\lambda^2}{8\pi^2 m_e^2 c^3 \epsilon_0} \int n_e \left[1 - 2 \frac{T_e}{m_e c^2}\right] \vec{B} \cdot \vec{dl} \quad (34)$$

$$\int n_e \left[1 - \frac{3}{2} \frac{T_e}{m_e c^2}\right] dl = \frac{4\pi c^2 m_e \epsilon_0}{e^2} \frac{\Delta\phi_{HeNe} \lambda_{HeNe} - \Delta\phi_{CO_2} \lambda_{CO_2}}{\lambda_{CO_2}^2 - \lambda_{HeNe}^2} \quad (35)$$

List of Publications

First Author Publications

A.J. Creely, T. Görler, G.D. Conway, S.J. Freethy, N.T. Howard, P.A. Schneider, A.E. White, M. Willensdorfer and The ASDEX Upgrade Team, “Validation study of GENE on ASDEX Upgrade using perturbative thermal diffusivity measured with partial sawtooth heat pulses,” *Nuclear Fusion* 58, 126001 (2018). <https://doi.org/10.1088/1741-4326/aadca1>.

A.J. Creely, S.J. Freethy, W.M. Burke, G.D. Conway, R. Leccacorvi, W.C. Parkin, D.R. Terry, and A.E. White, “Correlation electron cyclotron emission diagnostic and improved calculation of turbulent temperature fluctuation levels on ASDEX Upgrade,” *Review of Scientific Instruments* 89, 053503 (2018). <https://doi.org/10.1063/1.5005507>.

A.J. Creely, K. Ida, M. Yoshinuma, T. Tokuzawa, T. Tsujimura, T. Akiyama, R. Sakamoto, M. Emoto, K. Tanaka, and C.A. Michael, “Novel Analysis Technique for Measuring Edge Density Fluctuation Profiles with Reflectometry in the Large Helical Device,” *Review of Scientific Instruments* 88, 073509 (2017). <http://dx.doi.org/10.1063/1.4993437>.

A.J. Creely, N.T. Howard, P. Rodriguez-Fernandez, N. Cao, A.E. Hubbard, J.W. Hughes, J.E. Rice, A.E. White, J. Candy and G.M. Staebler, G.D. Conway, S.J. Freethy, and C. Sung, “Validation of Nonlinear Gyrokinetic Simulations of L- and I-mode Plasmas on Alcator C-Mod,” *Physics of Plasmas* 24, 056104 (2017). <http://dx.doi.org/10.1063/1.4977466>.

A.J. Creely, A.E. White, E.M. Edlund, N.T. Howard, and A.E. Hubbard, “Perturbative Thermal Diffusivity from Partial Sawtooth Crashes in Alcator C-Mod,” *Nuclear Fusion* 56, 036003 (2016). <http://dx.doi.org/10.1088/0029-5515/56/3/036003>.

Co-Author Publications

A.Q. Kuang, N.M. Cao, A.J. Creely, C.A. Dennett, J. Hecla, B. LaBombard, R.A. Tinguely, E.A. Tolman, H. Hoffman, M. Major, J. Ruiz Ruiz, D. Brunner, P. Grover, C. Laughman, B.N. Sorbom, and D.G. Whyte, “Conceptual design study for heat exhaust management in the ARC fusion pilot plant,” *Fusion Engineering and Design* 137, 221 (2018). <https://doi.org/10.1016/j.fusengdes.2018.09.007>.

P. Rodriguez-Fernandez, A.E. White, N.T. Howard, B.A. Grierson, G.M. Staebler, J.E. Rice, X. Yuan, N.M. Cao, A.J. Creely, M.J. Greenwald, A.E. Hubbard, J.W. Hughes, J.H. Irby, and F. Sciortino, “Explaining Cold-Pulse Dynamics in Tokamak Plasmas Using Local Turbulent Transport Models,”

Physical Review Letters 120, 075001 (2018). <https://doi.org/10.1103/PhysRevLett.120.075001>.

- S.J. Freethy, T. Görler, A.J. Creely, G.D. Conway, S.S. Denk, T. Happel, C. Koenen, P. Hennequin, A.E. White, and ASDEX Upgrade Team, “Validation of gyrokinetic simulations with measurements of electron temperature fluctuations and density-temperature phase angles on ASDEX Upgrade,” *Physics of Plasmas* 25, 055903 (2018). <https://doi.org/10.1063/1.5018930>.
- M. Schmidtmayr, J.W. Hughes, F. Ryter, E. Wolfrum, N. Cao, A.J. Creely, N. Howard, A.E. Hubbard, Y. Lin, M.L. Reinke, J.E. Rice, E.A. Tolman, S. Wukitch, Y. Ma, ASDEX Upgrade Team, and Alcator C-Mod Team, “Investigation of the critical edge ion heat flux for L-H transitions in Alcator C-Mod and its dependence on BT,” *Nuclear Fusion* 58, 056003 (2018). <https://doi.org/10.1088/1741-4326/aaaed0>.
- P. Rodriguez-Fernandez, A.E. White, A.J. Creely, M.J. Greenwald, N.T. Howard, F. Sciortino and J.C. Wright, “VITALS: A Surrogate-Based Optimization Framework for the Accelerated Validation of Plasma Transport Codes,” *Fusion Science and Technology* 74, 65 (2018). <https://doi.org/10.1080/15361055.2017.1396166>.
- N.T. Howard, C. Holland, A.E. White, M. Greenwald, P. Rodriguez-Fernandez, J. Candy, and A.J. Creely, “Multi-scale gyrokinetic simulations of an Alcator C-Mod, ELM-y H-mode plasma,” *Plasma Physics and Controlled Fusion* 60, 014034 (2018). <https://doi.org/10.1088/1361-6587/aa9148>.
- A.E. Hubbard, S.G. Baek, D. Brunner, A.J. Creely, I. Cziegler, E. Edlund, J.W. Hughes, B. LaBombard, Y. Lin, Z. Liu, E.S. Marmor, M.L. Reinke, J.E. Rice, B. Sorbom, C. Sung, J. Terry, C. Theiler, E.A. Tolman, J.R. Walk, A.E. White, D. Whyte, S.M. Wolfe, S. Wukitch, X.Q. Xu and the Alcator C-Mod Team, “Physics and performance of the I-mode regime over an expanded operating space on Alcator C-Mod,” *Nuclear Fusion* 57, 126039 (2017). <https://doi.org/10.1088/1741-4326/aa8570>.
- P. Rodriguez-Fernandez, J.E. Rice, N.M. Cao, A.J. Creely, N.T. Howard, A.E. Hubbard, J.H. Irby, and A.E. White, “On the correlation between ‘non-local’ effects and intrinsic rotation reversals in Alcator C-Mod,” *Nuclear Fusion* 57, 074001 (2017). <https://doi.org/10.1088/1741-4326/aa6e89>.
- N.T. Howard, C. Holland, A.E. White, M. Greenwald, J. Candy, and A.J. Creely, “Multi-scale gyrokinetic simulations: Comparison with experiment and implications for predicting turbulence and transport,” *Physics of Plasmas* 23, 056109 (2016). <http://dx.doi.org/10.1063/1.4946028>.
- S.J. Freethy, G.D. Conway, I. Classen, A.J. Creely, T. Happel, A. Köhn, B. Vanovac, and A.E. White, “Measurement of turbulent electron temperature fluctuations on the ASDEX Upgrade tokamak using correlated electron cyclotron

emission,” *Review of Scientific Instruments* 87, 11E102 (2016). <http://dx.doi.org/10.1063/1.4958908>.

A.E. White, N.T. Howard, A.J. Creely, M.A. Chilenski, M. Greenwald, A.E. Hubbard, J.W. Hughes, E. Marmor, J.E. Rice, J.M. Sierchio, C. Sung, J.R. Walk, D.G. Whyte, D.R. Mikkelsen, E.M. Edlund, C. Kung, C. Holland, J. Candy, C.C. Petty, M.L. Reinke, and C. Theiler, “Nonlinear gyrokinetic simulations of the I-mode high confinement regime and comparisons with experiment,” *Physics of Plasmas* 22, 056109 (2015). <http://dx.doi.org/10.1063/1.4921150>.

Chapter 1

Introduction

Unlimited Energy. For Everyone. Forever.
Fusion
It might actually work this time.

Lev Grossman
Time Magazine Cover
November 2, 2015

Fusion energy stands out among many possibilities as the most promising energy source to form the core of a sustainable energy future.

The benefits are almost too numerous to believe: no greenhouse gas emissions; no pollutants of any kind, for that matter; nearly unlimited fuel supply; no risk of meltdown; minimal radioactive material production; energy on demand; no geographical constraints and minimal land requirements. In other words, clean, safe energy wherever, whenever, and for as long as we need it.

And what's the catch? Well, fusion is hard. Very hard.

The problem boils down to almost exactly how fusion is described in popular culture references: making and sustaining an artificial star. Very generally, fusion involves combining small nuclei into larger, more stable ones, which releases energy in the process. Specifically, the most commonly attempted process involves fusing deuterium and tritium (hydrogen isotopes) into helium and an extra neutron. Each

nucleus, however, is positively charged, so in order to combine them, one must overcome the repulsive Coulomb force of two like charges. From the perspective of an individual particle, this means smashing the two nuclei together at high speed (the large kinetic energy is able to overcome the Coulomb repulsion). From a macroscopic perspective, having many nuclei all moving very quickly translates to generating very high temperatures. To overcome the Coulomb repulsion for some statistically meaningful amount of nuclei, temperatures on the order of one hundred million Kelvin are required. Anything at this temperature exists as a plasma: a collection of separated ions and electrons that is sometimes described as an ionized gas. The challenge of fusion energy is to create and sustain this hundred million degree plasma efficiently enough to generate net energy from the fusion reactions.

While there are several methods of achieving this goal, the most successful to date, and the topic of this thesis, is magnetic confinement. Since all of the particles in a plasma are charged, they all respond to the magnetic fields via the Lorentz force, orbiting around magnetic field lines. With a donut shaped magnetic field (known as a torus), particles can stream along the field without leaking out of an end of the machine, since there are no ends. Once one adds a vertical field and a current in the plasma to ensure that the plasma stays centered in the torus, one has what is known as a tokamak. This thesis will work with two tokamaks in particular, which are described in Chapter 2.

Even though the plasma is kept in place by the magnetic fields in the tokamak, heat (and particles) will slowly diffuse outward through collisions between particles (classical transport), which is accelerated due to larger particle orbits induced by the toroidal geometry of the machine (neoclassical transport). For this reason, one must continually heat the plasma in order to keep it hot enough for fusion (unless there is so much fusion power that the plasma heats itself, which is known as ignition). The ratio of fusion power generated in the plasma to heating power required to maintain

the plasma is known as the fusion gain, Q . Any fusion power plant must achieve a fusion gain of much greater than one (breakeven), likely somewhere around 10 or 20. Achieving the conditions necessary for breakeven has been the driving mission of fusion research for the last 60 years.

If the fusion plasma sat quietly in the tokamak, and one only had to worry about classical and neoclassical transport, a fusion power plant would probably have been built years ago. Unfortunately, when one tries to squeeze a hundred million degree plasma into a magnetic bottle, the plasma does anything but sit quietly: it roils and writhes and does its best to escape. More technically, the large temperature and density gradients in the plasma cause it to become unstable to various types of turbulence. This turbulence vastly increases the amount of heat that is transported outward compared to neoclassical levels, often by as much as a factor of 100. Understanding, and eventually predicting, this turbulence is one of the central challenges for fusion research, and is the focus of this thesis work.

One of the main motivations for understanding turbulence is so that one can use this knowledge in order to predict what will happen inside of a tokamak before it is built. An incredible amount of progress has been made on the path toward a fusion power plant since the first tokamaks were built, and until now, almost all of this innovation has been empirically derived: the design of new machines was informed by experimental results of previous machines, and improvements were based almost entirely on trends observed in experimental data. This process was certainly motivated by physical theories in conjunction with experimental data, but it has thus far not been feasible to design a machine based on first principles calculations. Truly complete simulations of turbulence are so computationally expensive as to be infeasible, and most simplified models to date have not included sufficient physics in order to accurately reproduce what is happening inside current machines, let alone predict what will happen in an entirely new device. The necessity of these models, however,

has become increasingly clear, as state-of-the-art fusion experiments have become increasingly expensive and slower to construct. Reliable predictions of turbulence and its effects on performance in fusion plasmas could lead to significant improvements in future fusion experiments and eventually fusion power plants.

Fortunately, immense progress has been made in the last few decades in our ability to measure the effects of turbulence in tokamak plasmas, our theoretical understanding of turbulence, and in the ability of turbulent transport simulations to accurately model the transport of heat and particles inside of fusion experiments. These three elements are complementary: measurements of turbulence provide additional information that aid the advancement of turbulence theory, theory is implemented in new simulations, simulations reveal new aspects or implications of theory, and both theory and simulation motivate new measurements in order to confirm results. This thesis focuses primarily on experimental measurements of turbulence and its effects, and comparison of these measurements to the results of turbulent transport simulations. This inevitably leads to some discussion of the physics of turbulence and its role in transport, though only limited attention is dedicated to that topic.

Before delving into the experimental and modeling contributions of this thesis, Chapter 2 goes over the background information and general concepts that are key to understanding the rest of the work performed in this thesis. In particular, a very general introduction to turbulence in tokamaks is required, including the basics of why turbulence develops in tokamak plasmas, how this leads to transport of heat and particles, and what turbulent modes are dominantly present in high performance tokamaks. In addition to general turbulence theory, Chapter 2 introduces a few models for turbulent transport, specifically the gyrokinetic and gyrofluid models, as well as the specific code implementations of both that are used in this thesis (GYRO, GENE, and TGLF). One of the important questions raised when discussing these turbulent transport models is whether ion-scale simulations sufficiently capture the relevant

plasma phenomena to predict future machine performance, or whether one requires multi-scale simulations. The chapter then goes on to describe the two tokamaks studied in this thesis, Alcator C-Mod and ASDEX Upgrade. Finally, this background chapter describes the process of validation, and introduces the specific validation constraints used in this thesis (electron and ion heat fluxes, electron perturbative thermal diffusivity, and electron temperature fluctuations).

The contributions and results of this thesis are described in Chapters 3 through 6, with Chapters 3 and 4 describing hardware and data analysis contributions intended on measuring validation constraints experimentally, and Chapters 5 and 6 using these measurements in validation studies of various turbulent transport models.

Chapter 3 focuses on one of the validation constraints mentioned in Chapter 2, the perturbative thermal diffusivity. The perturbative thermal diffusivity measures how sensitive the turbulent heat flux in the plasma is to the driving temperature gradient. This thesis developed a new method of measuring the perturbative diffusivity experimentally, based on tracking heat pulses generated by partial sawtooth crashes. Chapter 3 describes the physical principles behind this method, outlines the required calculations, shows comparisons with previous methods of measuring the same quantity, and presents experimental observations from Alcator C-Mod and ASDEX Upgrade. This chapter focuses on developing the measurement technique and demonstrating its readiness for use as a validation constraint, though some purely experimental results are also included.

Chapter 4 focuses on a second validation constraint that is important to this thesis, electron temperature fluctuations. In particular, two correlation electron cyclotron emission (CECE) diagnostic systems were constructed as part of this thesis, one on Alcator C-Mod and one on ASDEX Upgrade. The chapter begins by describing the general operation of ECE radiometers, before describing in detail the hardware that was constructed for and installed on both machines. In addition to

hardware, this thesis work also contributed newly derived analysis techniques in order to extract temperature fluctuation levels from raw CECE data. These calculation techniques are both outlined and tested in Chapter 4. Finally, the chapter presents example experimental data from both machines, as well as some interesting experimental observations that stand separate from the validation comparisons presented in subsequent chapters.

Using these two validation constraints, as well as the more commonly used electron and ion heat fluxes, Chapter 5 presents the results of four validation studies of nonlinear gyrokinetic codes: two with GYRO on Alcator C-Mod and two with GENE on ASDEX Upgrade. The first validation study of GYRO on Alcator C-Mod looks at an L- and an I-mode plasma and applies all four validation constraints, finding rather poor agreement between ion-scale GYRO simulations and the experimental measurements. The second study, which involved adding the perturbative diffusivity to previously performed validation efforts, showed that in at least some cases, the experimentally measured perturbative diffusivity disagreed with ion-scale GYRO simulations but agreed with multi-scale GYRO simulations. Moving on to validation of GENE on ASDEX Upgrade, two studies at different radial locations of the same plasma discharge are presented. The first, which was performed entirely as part of this thesis, used heat fluxes and perturbative diffusivity as constraints, and found good agreement between ion-scale GENE simulations and all three constraints. The second, which included significant contributions from this thesis but was performed separately, included heat fluxes and electron temperature fluctuations as constraints (as well as two others that are briefly discussed), finding good agreement with four out of five constraints and ion-scale GENE simulations. This chapter also includes some preliminary suggestions as to how one might distinguish plasmas in which multi-scale physics are important from those for which ion-scale simulations are sufficient.

Chapter 6 more directly addresses the question of the importance of multi-scale

physics in turbulent transport simulations. In order to compare directly ion- and multi-scale simulations without allocating enormous computational resources, this chapter utilizes a quasi-linear gyrofluid model, TGLF. This new model is first applied to the same L- and I-mode plasmas studied in the first validation study in Chapter 5, showing that merely including multi-scale effects did not resolve the discrepancies between simulation and experimental constraints. This result reveals that another explanation for discrepancies between simulation and experiment may be uncertainties in inputs beyond just electron and ion temperature gradients. The final validation study in this thesis builds upon the results of all of the other work in this thesis. This study validated TGLF in both ion- and multi-scale configurations, varying multiple inputs and validating against up to four constraints, on 11 plasma discharges on Alcator C-Mod and ASDEX Upgrade. In addition to showing that multi-scale TGLF agrees with all available validation constraints on all of the discharges in the study, the results of this study are used to develop two criteria with which to predict whether or not multi-scale turbulent effects are important in a given plasma. Both criteria are based on the relative linear growth rates of different turbulent modes, and provide some insight into the fundamental physical nature of multi-scale turbulence, and how the rather disparate scales interact with one another.

Finally, Chapter 7 summarizes the results of this thesis, and discusses how they fit into both the larger turbulent transport validation effort and the overarching march toward a fusion power plant. While the work of this thesis is but one aspect of a much larger endeavor, spanning continents and generations, it represents a concrete step on the path to fusion energy.

Chapter 2

Background and General Concepts

This chapter briefly introduces many of the key concepts and tools that are used throughout this thesis. First, some central concepts of turbulence and transport in tokamak plasmas, which are most relevant to this thesis work, are outlined, including references for further reading. This thesis makes use of three turbulent transport simulation codes, GYRO, GENE, and TGLF, each of which is briefly described. The two tokamaks studied in this thesis, Alcator C-Mod and ASDEX Upgrade, are then both described, along with the more standard diagnostics used to measure the background plasma equilibrium in these machines. This chapter then describes the concept of validation, which is how one checks that a turbulent transport model is correctly modeling the relevant properties of the plasma turbulence. Finally, this chapter describes the experimental measurements that are used as constraints in this validation process, including the two that were the primary focus of this thesis, electron temperature fluctuations and electron perturbative thermal diffusivity.

2.1 Turbulence and Transport in Tokamaks

This section introduces the concept of drift wave turbulence and how it relates to the transport of heat and particles in fusion plasmas. While turbulence theory is

an incredibly rich field of research, and is intimately linked to both experimental measurements and simulations of tokamak plasmas, the theory itself is not the primary focus of this thesis, and so it is only explained in a very introductory manner, sufficient to understand the results presented in the main body of this thesis. In addition, this section focuses solely on turbulence in tokamak plasmas, though there are many other forms of plasma turbulence in other situations, and also many other types of waves besides drift waves that are supported in tokamak plasmas. There are many more complete reviews of turbulent transport in tokamaks in the literature, and the interested reader is directed to References [2, 3, 4, 5], among others. Reference [6] also contains a wealth of general plasma physics background, including both fundamental physics and engineering considerations regarding tokamaks and other magnetically confined fusion concepts. The explanation in this section will largely follow References [5, 7], discussing turbulence in general, a few specific turbulent modes in tokamaks, and the concept of saturation.

In general, turbulence is characterized by chaotic fluid motion driven by gradients in fluid properties (or by shearing and other asymmetries) and saturated by nonlinear interactions of phenomena of different scales. In tokamaks, the fluid in question is the plasma, which is generally formed out of hydrogen isotope ions (typically deuterium) and electrons (as well as small amounts of heavier ions, generally referred to as impurities). The plasma is assumed to be quasi-neutral, meaning that the plasma does not have a net charge. In heating the plasma to the necessary hundred million Kelvin required to fuse hydrogen to helium, and achieving sufficient plasma density to fuse at a reasonable rate, one inevitably generates gradients in both temperature and density, with large temperatures and density in the core of the tokamak torus, and low temperatures and densities at the edge of the plasma. It is these gradients that provide the free energy to drive turbulence unstable, which eventually leads to transport of heat and the relaxation of these gradients (unless, of course, a source of

heat exists, such as heating from fusion reactions or external heating of the plasma).

In particular, this thesis focuses on core turbulence. Tokamak plasmas are generally separated into several radial regions, depending on the dominant transport mechanisms and the models used to describe each region. Different literature may use slightly different region delineation, but this thesis will define four regions. Starting from the outside, the first region is the scrape off layer (SOL). The SOL is the region of open magnetic field lines, which do not loop back on themselves. This region therefore contains unconfined plasma, and typically has quite low temperatures. The inner boundary of the SOL is the separatrix, beyond which one finds the edge and pedestal region. The edge and pedestal region is characterized by steep gradients and coupling to the SOL, and extends from the separatrix in to roughly $\rho_{tor} \approx 0.8 - 0.9$, where ρ_{tor} is the square root of the normalized toroidal magnetic flux.¹ Next, is the core turbulence region, which extends from the top of the pedestal in to the sawtooth mixing radius. Transport in this region is dominated by the types of turbulence described in Section 2.1. Finally, inside of the sawtooth mixing radius, transport caused by macroscopic plasma motion dominates, and turbulent transport is less important. The sawtooth instability, which is discussed in some further detail later in this thesis, is caused by a magnetohydrodynamic (MHD) kink instability, and leads to large scale mixing of the plasma inside of a surface that depends on the relative values of the toroidal and poloidal magnetic fields [6]. Note that if the plasma current is low enough, or if the current profile is sufficiently flat, then the plasma may not sawtooth at all. With these four regions laid out, consideration returns now to just core turbulence.

¹Normalized radial coordinates are common in tokamak research. The three most common are ρ_{tor} , ρ_{pol} , and r/a . All three range from 0 to 1, with 0 as the center of the plasma and 1 as the separatrix. ρ_{tor} and ρ_{pol} are both normalized based on magnetic flux, with ρ_{tor} based on the toroidal magnetic field and ρ_{pol} based on the poloidal magnetic field. The last, r/a is normalized based on physical distance. If the magnetic geometry within the tokamak is known, it is possible to convert between the three normalizations. Different normalizations are more or less convenient for different calculations, but it is important to be clear which normalization one is using. This thesis will almost exclusively make use of ρ_{tor} .

Most turbulence in the tokamak core is driven unstable by what is known as drift waves. Fairly general plasma physics calculations show that in a magnetically confined plasma, gradients in ion and electron pressure (either density or temperature) lead to diamagnetic currents in the plasma, which are necessary for macroscopic equilibrium. If one introduces a perturbation to the pressure gradient, then a corresponding perturbation to the diamagnetic current develops, which propagates perpendicular to both the background magnetic field and the pressure gradient. In order to maintain quasi-neutrality, a parallel current also forms, mostly carried by the electron due to their small mass. Waves propagating in the ion diamagnetic direction are known as ion drift waves, and similarly waves propagating the electron diamagnetic drift direction are known as electron drift waves.

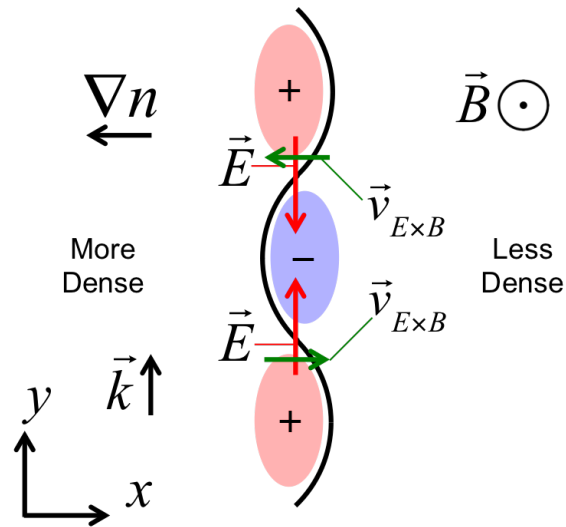


Figure 2-1: An illustration of an ion density perturbation propagating as a drift wave in a background ion density gradient and magnetic field. Modified from Reference [8].

As an example, consider an ion density perturbation, illustrated in Figure 2-1. Regions of higher density develop a slightly higher potential, and regions of lower density a slightly lower potential. This then generates an electric field perpendicular to the gradient. The electric field then causes ions to drift at their $E \times B$ drift velocity either along the gradient or against the gradient, depending on the sign of

the electric field at each location in the wave. This motion serves to propagate the density perturbation perpendicular to both the background magnetic field and the density gradient. A very similar argument can be made for temperature gradients.

If there is no dissipation of the electron motion, then the density and potential perturbations are in phase, and there is no net transport caused by the wave: the wave stably propagates without affecting the background gradients. If, on the other hand, there is some dissipation of the parallel electron motion (caused by electron-ion collisions, wave particle interactions, or interactions with particles mirror-trapped in the spatially varying toroidal magnetic field), then there will be a phase delay between the density and potential perturbations. This delay generates a linear instability, where the drift wave grows in amplitude. This growing wave leads to transport of particles down the gradient. Note that the phase delay between density and potential can also arise from ion dynamics, even with adiabatic electrons. Again, a similar process occurs for temperature gradients.

These drift wave perturbations initially grow linearly, but eventually reach sufficient amplitude that they begin to interact nonlinearly with phenomena at other wavenumbers (where the wavenumber is 2π divided by the perturbation wavelength). This then leads to perturbations at these other wavenumbers, some of which will be linearly stable and thus damped by various mechanisms within the plasma. These new waves then couple to yet more wavenumbers and produce a broad spectrum of turbulence wavenumbers and frequencies. The continuous injection of energy into some wavenumbers, nonlinear coupling to other wavenumbers, and eventual dissipation leads to a saturated state of the turbulence, in which the amplitude of the turbulence remains constant in time (assuming the background plasma equilibrium remains the same).

The specific nonlinear interactions of the drift waves depend on how they are driven unstable, and different types of turbulence have different characteristic density,

temperature, and potential fluctuations, as well as different cascades of energy to larger and smaller scales. One major defining characteristic is the overall spatial scale on which the turbulent features exist. This is typically expressed in terms of the ion gyro-radius evaluated at the electron temperature, $\rho_s = c_s/\omega_{ci}$, where $c_s = \sqrt{2T_e/m_i}$ is the ion sound speed and $\omega_{ci} = q_i B/m_i$ is the ion gyro-frequency. In these equations, T_e is the electron temperature, m_i is the ion mass, q_i is the ion charge, and B is the magnetic field. For example, ‘ion-scale’ turbulent features typically exist on the scale of ρ_s .

While different relative phases of the temperature, density, and potential fluctuations can drive heat, particles, and momentum in the plasma, this thesis focuses on heat transport in particular. This heat transport contains both electrostatic (only the electric potential varies) and electromagnetic (the magnetic field also varies) contributions.

Different turbulent modes are typically categorized by the dominant driving gradient or source of instability. While there are many turbulent modes in different plasma conditions and geometries, the three dominant modes in most tokamak plasmas, and thus the three of most interest in this thesis, are the following: the ion temperature gradient (ITG) mode, the trapped electron mode (TEM), and the electron temperature gradient (ETG) mode. While detailed derivations of each of these modes is omitted for the purpose of brevity (see Reference [4] for such work), it is worthwhile to outline the general characteristics of each mode that are most relevant for the work in this thesis.

The ion temperature gradient (ITG) mode is, perhaps unsurprisingly, driven unstable by gradients in the ion temperature. Specifically, one is interested in the normalized gradient, often expressed in terms of the normalized gradient scale length, a/L_{T_i} , where $L_{T_i} = T_i(\partial T_i/\partial r)^{-1}$ is the gradient scale length and a is the plasma minor radius. Note that some literature normalizes by the plasma major radius, R ,

rather than the minor radius, but the essential physics is the same. The ITG generally has some critical value of a/L_{Ti} , below which the mode is stable and above which the mode is driven unstable, though the exact value of the critical gradient depends on details of the geometry and background plasma. The turbulent structures resulting from the ITG instability typically exist on scales of $k_\theta \rho_s \approx 0.1 - 1.0$, where k_θ is the poloidal wavenumber of the turbulence and ρ_s was defined above. This range of wavenumbers is commonly referred to as the ‘ion-scale’ range, since it encompasses phenomena larger than the ion gyro-radius. This turbulence propagates in the ion diamagnetic drift direction, and can transport heat through both the ion and electron channels.

Another ion-scale turbulent mode is the trapped electron mode (TEM). The TEM is a result of trapped electrons (trapped in the low-field side of the magnetic flux surface, as if in a magnetic mirror, due to the radial variation of the toroidal magnetic field) causing a phase delay in the drift wave. This mode is driven primarily by electron pressure gradients (either electron density or temperature), and its strength depends on the fraction of trapped electrons. For this reason, factors that affect the fraction of trapped electrons, such as the aspect ratio and the collisionality, also influence the stability of the TEM. The TEM also generally has a critical pressure gradient, expressed as the density (a/L_{ne}) and electron temperature (a/L_{Te}) normalized gradient scale lengths. TEM turbulence propagates in the electron diamagnetic direction, and exists up to wavenumbers slightly above $k_\theta \rho_s \approx 1.0$.

At electron scales, an analogue to the ITG exists, known as the electron temperature gradient (ETG) mode. The ETG is very similar to the ITG in many respects, except that it operates on the electron scales, $k_\theta \rho_s \approx 1.0 - 60.0$. There is often a continuous transition between the TEM and ETG at intermediate wavenumbers, such that there is no clear boundary between the two types of modes. The ETG is sensitive to the electron temperature gradient, a/L_{Te} , propagates in the electron diamagnetic

direction, and drives electron heat flux. Though the ETG exists on electron scales, and thus one might think that the small eddy size compared to ion scale turbulence would lead to negligible transport, it can still be responsible for considerable heat transport due to the existence of radially elongated streamers, which significantly increase the effective eddy size [9].

In addition to these three dominant turbulent modes in tokamak plasmas, a phenomenon known as zonal flows also plays a key role in determining the final turbulent state of the plasma (in addition to the nonlinear transfer of energy between scales discussed above). While the theoretical description of how zonal flows are generated quickly becomes quite mathematically complicated, qualitatively they result from small scale turbulent velocity fluctuations forming larger structures with ordered flows perpendicular to the background driving gradients. Zonal flows have radial structure on the order of the background density gradient scale length (often a few centimeters), and are very poloidally elongated (approximately poloidally symmetric). Zonal flows do not create or dissipate momentum, they simply rearrange it. They are linearly stable and do not induce transport. Instead, their variation with radius and in time tends to break up turbulent eddies, reducing the turbulent correlation length and correspondingly reducing turbulent transport. To summarize, turbulence gives rise to zonal flows, which then act to dampen the turbulence. The interaction between turbulence and zonal flows is therefore key to understanding the final saturated state of the plasma turbulence.

In particular, zonal flows are one of the key mechanisms through which ion- and electron-scale turbulence interact [10]. In general, turbulence is dominated by a cascade of energy from large scales down to small scales. This is generically how energy from large scale gradients is transferred to smaller and smaller eddies, with each scale being fed energy by the adjacent larger scale and feeding energy to the adjacent smaller scale, until at some point the energy dissipates (via, for example, viscos-

ity in normal fluid mechanics). As will be discussed in greater detail in Section 2.2, however, recent simulations of plasma turbulence have revealed that ion- and electron-scale turbulence interacts strongly, and electron-scale features, such as the ETG, have been shown to affect ion-scale turbulence [10, 11, 12, 13, 14, 15, 16, 17]. Some of this interaction appears to be due to direct inverse cascades or non-local (in wavenumber space) energy transfer, but zonal flows also seem to play an important role in these interactions. Specifically, both ion- and electron-scale turbulence influences zonal flows, which then affect both ion- and electron-scales. As an example, increased electron-scale turbulence has been shown to reduce the amplitude of zonal flows generated by ion-scale turbulence, thereby increasing the saturation level of the ion-scale turbulence. For this reason, zonal flows play a key role in both single-scale and multi-scale turbulence saturation.

While this section has only scratched the surface of an incredibly rich field of turbulence theory, it has hopefully provided the introduction necessary to understand the results presented in the remainder of this thesis. With this background in mind, this thesis will now introduce a few specific examples of turbulent transport models that are utilized in this work, before moving onto experimental hardware and the process of using experiment to validate models.

2.2 Turbulent Transport Models

In order to better understand the complicated nature of turbulence in fusion plasmas, and work toward using this knowledge to predict the performance of future fusion devices, a number of models have been developed through the years that describe plasma turbulence in varying levels of completeness. Just as with turbulence theory, this section does not hope to provide a detailed description of turbulence modeling in general or of any of the specific implementations used in this thesis. Instead, this section hopes to provide the reader enough background to understand the motivation

behind and results of this thesis work.

While the fusion community has developed numerous models to simulate turbulent transport in the core of tokamaks, only two of the most common are described and used in this thesis: the gyrokinetic and gyrofluid models. Both of these models are quite complex, since turbulence is a fundamentally complicated phenomenon, and so practical uses of both models involve implementations in numerical simulations. In particular, this thesis works with two gyrokinetic simulation codes, GYRO [18] and GENE (Gyrokinetic Electromagnetic Numerical Experiment) [9], and one gyrofluid code, TGLF (Trapped Gyro-Landau Fluid) [19, 20, 21, 22]. The remainder of this section discusses the basics of the gyrokinetic and gyrofluid models (largely following References [23, 24]), as well as touching briefly on these three specific implementations.

In a world unconstrained by computation time, it is in theory possible to fully describe nearly all behavior of any plasma system by coupling Newtonian mechanics with Maxwell's equations. The plasma is made up of many charged particles which interact primarily through long range collective electromagnetic effects and binary Coulomb collisions. True, there are some very specific situations in which quantum effects become important (i.e. Motional Stark Effect diagnostics), and the fusion event itself is fundamentally not a classical process, but on the whole, plasmas are mostly governed by purely classical physics. The immediate issue with this approach is that the plasma contains of order 10^{20} or more particles, making any direct computation of each particle's motion completely infeasible in the foreseeable future.

For this reason, nearly all models of plasma turbulence utilize some statistical approach. Generally one begins with a kinetic formulation, using the Boltzmann Equation (or the Vlasov Equation in the collisionless case) coupled to Maxwell's Equations. The Boltzmann Equation can be expressed as:

$$\frac{\partial f_s}{\partial t} + \vec{v} \cdot \nabla f_s + \vec{a} \cdot \nabla_v f_s = C(f_s, f_{s'}) \quad (2.1)$$

where $f_s(\vec{x}, \vec{v}, t)$ is the distribution function (the density in phase space) of a given species in the plasma, \vec{a} is the acceleration due to external forces (i.e. electromagnetic fields) and $C(f_s, f_{s'})$ is a collision operator, accounting for the effects of binary collisions between particles in the plasma.

This is then coupled to Maxwell's Equations:

$$\nabla \cdot \vec{E} = \frac{1}{\epsilon_o} \sum_s q_s n_s \quad (2.2a)$$

$$\nabla \cdot \vec{B} = 0 \quad (2.2b)$$

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \quad (2.2c)$$

$$\nabla \times \vec{B} = \mu_o \sum_s \vec{J}_s + \mu_o \epsilon_o \frac{\partial \vec{E}}{\partial t} \quad (2.2d)$$

where \vec{E} is the electric field, \vec{B} is the magnetic field, q_s is the charge of species s , n_s is the number density of species s , and \vec{J}_s is the current density of species s , defined as:

$$n_s = \int f_s d\vec{v} \quad (2.3)$$

and

$$\vec{J}_s = q_s \int \vec{v} f_s d\vec{v} \quad (2.4)$$

Collectively, this set of equations is known as the kinetic treatment.

Even though treating the plasma statistically reduces the computational requirements significantly, the set of the Boltzmann Equation and Maxwell's Equations are

still incredibly complex when applied to any realistic plasma system. There are still six dimensions in the system (three spatial dimensions and three velocity dimensions) and relevant plasma phenomena span many orders of magnitude of spatial and temporal scales. For example, the plasma system may be meters in size, while the electron gyro-radius may be on the order of 10^{-4} meters. In time, macroscopic temperature and density profile evolution may occur on the order of seconds, while the electron cyclotron frequency may be on the order of 10^{-11} seconds. Even if one is only concerned with one radial location in the core plasma in steady macroscopic equilibrium, relevant phenomena still span many orders of magnitude. Finally, turbulence is a highly nonlinear phenomenon, and the background magnetic field and toroidal geometry make the plasma dynamics highly anisotropic.

The Gyrokinetic Model

The gyrokinetic model makes further simplifications to the kinetic model, taking advantage of separations in relevant scales. Most simply, the gyrokinetic model reduces the number of dimensions in the kinetic equations from six to five, by taking advantage of the small size of the gyro-radius compared to other scales in the system. Instead of separately keeping track of x, y and z velocities, one keeps track of only two velocities, parallel and perpendicular to the magnetic field. In line with this, the model tracks the guiding center of the particle, the point around which the particle executes its gyro-orbit, instead of tracking the exact particle position during its orbit.

More mathematically, the gyrokinetic model is derived with the following ordering in the small parameter ϵ_{GK} [23]:

$$\frac{\omega}{\omega_{cs}} \sim \frac{k_{\parallel}}{k_{\perp}} \sim \frac{\tilde{v}_E}{v_{Ts}} \sim \frac{\tilde{n}_s}{n_o} \sim \frac{\tilde{B}_1}{B_o} \sim \frac{\rho_{Ls}}{L_{ns}} \sim \mathcal{O}(\epsilon_{GK}) \quad (2.5)$$

where ω is characteristic turbulent frequency, ω_{cs} is the cyclotron frequency of species s , k_{\parallel} and k_{\perp} are the turbulence wavenumbers parallel and perpendicular to

the background magnetic field, \tilde{v}_E is perturbed $E \times B$ drift velocity, v_{Ts} is the thermal velocity of species s , \tilde{n}_s is the perturbed density, n_o is equilibrium density, \tilde{B}_1 is the perturbed magnetic field, B_o is the background magnetic field, ρ_{Ls} is the gyro-radius of species s , and L_{n_s} is density gradient scale length of species s (as defined above).

Since this thesis does not focus on the theoretical aspects of the gyrokinetic model, the full forms of the final set of gyrokinetic equations are not repeated here, and can be found in, for instance, Reference [23]. Instead, more attention is given to the various additional simplifications that are generally made when one actually runs gyrokinetic simulations of tokamak plasmas, and the impact that these have on the fidelity of the results. A number of differences are briefly touched upon, though the last, the differences between ion-, electron-, multi-scale simulations, is the most relevant to the work presented in this thesis. This list is not meant to be exhaustive, and is primarily intended to illustrate that even within the gyrokinetic model, many differences exist in how one implements the model into simulations.

Numerical Approach

First, there are three main numerical approaches that have been used to solve the set of gyrokinetic equations: the Lagrangian method (often called particle-in-cell, based on sampling initial positions in phase space, moving particles based on forces, and then recalculating the fields at each timestep), the Eulerian method (using a fixed grid in phase space and finite difference to calculate the distribution function at the next time step before recalculating field equations), and the semi-Lagrangian method (using a fixed grid and integrating orbits back in time, then interpolating to the next time step and recalculating fields) [23]. Each method has various computational advantages and disadvantages. The details of these numerical approaches are not the focus of this work, though it is worth noting that both gyrokinetic codes used in this thesis, GYRO and GENE, are Eulerian codes. Codes are also often categorized by whether they

compute the full distribution function (a ‘full-f’ code) or approximate the distribution function as Maxwellian and then compute the deviation from Maxwellian (a ‘delta-f’ code).

Collisions

Another important consideration is the exact formulation of the collision operator found on the right hand side of Equation 2.1. Various codes will parameterize this operator differently, with some setting this to zero since collisional time scales are typically much longer than turbulence time scales (in the collisionless case, the Boltzmann Equation is generally known as the Vlasov Equation). This is one of the main differences, for example, between GYRO and the more recently released CGYRO [25]. One example of a collision operator is the linearized Landau-Boltzmann operator, which is used in the GENE code [26].

Global and Local Simulations

Considering now approximations and simplifications that can often be turned on or off within a given code, one of the most common choices is whether one simulates the entire core turbulence region (a ‘global’ simulation) or only one radial location (a ‘local’ or ‘flux-tube’ simulation). Global codes simulate a large radial domain and take into account a large portion of the plasma geometry, allowing all equilibrium plasma parameters to vary across the domain. In theory, these codes can be allowed to self-consistently evolve profiles and can even be used to simulate time-dependent phenomena. Unsurprisingly, the main drawback of global simulations is the significantly larger computational expense. In contrast, local codes only simulate the plasma in a narrow ‘flux-tube,’ which is the region in the vicinity of a single magnetic field line, using periodic boundary conditions. These simulations use fixed background gradients, and cannot simulate shear in the background gradients. The

advantage is that local simulations require significantly less computational time than global simulations.

Electrostatic and Electromagnetic Simulations

One of the more common binary settings in gyrokinetic simulations is whether or not to include the effects of electromagnetic turbulence. If one goes through the gyrokinetic equations in detail, one finds that the contributions of the varying magnetic field to the overall properties of the turbulence depend strongly on the plasma $\beta = p/(B^2/2\mu_o)$, which is the ratio of the plasma kinetic pressure to the magnetic pressure. If β is sufficiently low, the effect of the magnetic fluctuations is negligible, and simulations can leave these terms out completely, in what is known as an ‘electrostatic’ simulation. If, on the other hand, the plasma β is large (on the order of 0.01), then contributions from these terms become important, and one runs what is known as an ‘electromagnetic’ simulation. Some codes, such as GENE, default to the electromagnetic case, while others, such as GYRO, can be run in either configuration.

Drift-Kinetic and Adiabatic Electrons

Another simplification that is fairly commonplace is to model ions using the gyrokinetic model, but to use some further simplified model for the electrons. For example, electrons may be treated with the drift-kinetic model ($k_\theta \rho_e \rightarrow 0$) or may be treated adiabatically [23]. These simplifications were motivated by theories suggesting that ions would dominate heat transport in tokamak plasmas due to their larger gyro-radius. Practically, these simplifications considerably reduced the required computational expense, which was absolutely vital in the earlier days of plasma simulations when supercomputers were slower. The adiabatic approximation in particular and the electron drift-kinetic approximation to a lesser degree, however, were shown to leave out important phenomena even on the ion scales [23]. In addition, some more

recent work has shown that electron contributions to the overall turbulent state of the plasma may be quite important, and that one may therefore be required to use fully gyrokinetic electrons [10, 11, 12, 13, 14, 15, 16, 17]. This is discussed in greater detail in the next paragraph.

Ion-, Electron-, and Multi-Scale Simulations

Finally, one of the most important choices that one makes when running a gyrokinetic simulation is the resolution in spatial and time scales. Intuitively, if one uses a grid size that only resolves down to the ion gyro-radius, anything occurring on the scale of the electron gyro-radius will be washed out. Similarly, if the entire simulation domain is smaller than an ion gyro-radius, than any large scale phenomena are left off. Historically, most gyrokinetic simulations have been run in an ‘ion-scale’ configuration, which resolves scales down to approximately the ion gyro-radius ($k_\theta \rho_s \approx 1.0$ in wavenumber space). This was primarily to save computational time, and was motivated by the theory just discussed, that ions were the dominant source of heat transport in plasmas. In some cases, however, it was discovered that the ETG mode contributed significantly to heat transport, and researchers responded by also simulating just the ETG, in what is known as an ‘electron-scale’ simulation, resolving $k_\theta \rho_s \approx 2.0 - 60.0$.

Only within the last five years or so has computational power advanced to the point where one can run simulations that simultaneously resolve electron- and ion-scales ($k_\theta \rho_s \approx 0.1 - 60.0$), in what is known as multi-scale simulations [10, 11, 12, 13, 14, 15, 16, 17]. These simulations are enormously computationally expensive (with a single study consuming as much as 10^8 CPU-hours), but have provided incredible insight into the manner in which ion- and electron-scales interact nonlinearly. These simulations have revealed that while in some cases ion-scale simulations are sufficient to model the plasma dynamics, in others multi-scale simulations are absolutely

required to accurately capture the turbulent state. The process of comparing the results of these simulations to experimental measurements, validation, is discussed in greater detail later in Section 2.4, but one of the main goals of this thesis is to use the validation process to determine when multi-scale effects are important, and when ion-scale simulations sufficiently model the plasma turbulence. This topic is discussed in greater detail throughout this thesis.

Other Parameters

In addition to the many computational choices and simplifications listed here, there are many others which are discussed throughout the literature. To name just a few, there are many different ways to represent the geometry of the tokamak plasma, with the Miller parameterization being one of the most common [27]. Impurities (other ion species besides deuterium, or whatever the dominant species is) may also play an important role in determining the fully turbulent state [23], and one can either include a third gyrokinetic impurity species, or attempt to model the effects of the impurity through an effective charge of the plasma Z_{eff} . Fast ions, generated by neutral beams, ion cyclotron heating, fusion products, or some other source, may also have an effect on the background turbulence [28, 29]. There are many other modeling choices that are not discussed here, and Reference [23] gives a more thorough review.

Recall that the two gyrokinetic models used in this thesis were GENE and GYRO. A list of the code settings chosen for each simulation presented in this thesis will be presented along with the results, as different simulations in this thesis used different settings. Briefly, however, a quick description of GENE and GYRO is warranted here. Both codes use the Eulerian numerical approach to solve the gyrokinetic equations and both are delta-f codes. GENE is inherently electromagnetic, while GYRO can run in either electrostatic or electromagnetic configurations. The codes use different collision operators, described in their respective literatures. Both codes can use gyrokinetic,

drift-kinetic, or adiabatic electrons, and can be run in ion-, electron-, or multi-scale configurations. Finally, both codes can be run linearly, returning linear growth rates (the growth rate of a single turbulent mode at a particular wavenumber, uncoupled to other modes) and real frequencies (the frequency of a given unstable mode), and nonlinearly, returning the saturated turbulent state and all of its associated properties (heat fluxes, fluctuations levels, etc.). Both GENE and GYRO can also run either local or global simulations.

This completes the brief description of the gyrokinetic model and the various choices one is faced when implementing the model into a simulation. As mentioned at the beginning of this section, however, this thesis also makes use of a quasi-linear gyrofluid model, TGLF, which also merits a brief qualitative description.

The Gyrofluid Model

Just as one can take moments of the kinetic equations in order to develop a fluid picture of liquids and gasses, one can take moments of the gyrokinetic system of equations to develop a fluid picture of a magnetized plasma. A model based on these approximations is known as a gyrofluid model [19]. As with all fluid descriptions, one must close the model somehow, and if one does this in such a way as to retain the most relevant kinetic effects, such as Landau damping, the model is known as a Gyro-Landau Fluid. The TGLF code also includes the effects of trapped particles in its fluid model, hence the name Trapped Gyro Landau Fluid. TGLF uses this fluid model in order to calculate the linear growth rates of turbulent modes, and has been shown to have good agreement with linear growth rate calculates from gyrokinetic simulations over a wide parameter range [19, 20, 21, 22]. Since the gyrofluid picture is considerably simpler than the gyrokinetic picture, these calculations are also considerably faster.

Another significant difference between TGLF and GYRO or GENE is that most gyrofluid models do not describe the nonlinear saturated state of plasma turbulence.

Instead, these models are known as ‘quasi-linear.’ TGLF calculates the linear growth rates of the turbulence, and then uses a built-in saturation rule to translate these linear growth rates to a saturated plasma state. It does not directly calculate the nonlinear interactions. The saturation rule is therefore very important, and in the case of TGLF, it was originally generated by fitting a model to the results of several hundred single-scale nonlinear GYRO simulations [20]. The original saturation rule, known as SAT0, was then improved by adding multi-scale GYRO simulations to the model, which resulted in the current saturation rule, known as SAT1 [21]. In addition to TGLF, another commonly used quasi-linear model is QuaLiKiz, which uses the results of nonlinear GENE simulations to build its saturation rule (note that QuaLiKiz uses a gyrokinetic model to calculate growth rates, rather than a gyrofluid model) [30].

TGLF has both advantages and disadvantages when compared to a nonlinear gyrokinetic code such as GYRO or GENE, a few of which are listed here. First, TGLF is able to calculate many of the same quantities, such as heat fluxes and fluctuation levels, and do so much more quickly. For example, TGLF is able to run multi-scale simulations in a matter of seconds, many orders of magnitude faster than a nonlinear gyrokinetic code. This greatly expands the possibilities for predictive work and large scale validation studies, which is the main way in which the code will be used in this thesis. On the other hand, TGLF cannot be run as a global code, and instead uses a series of local runs in order to build up a radial profile when used for that purpose. Perhaps most significantly, the fact that TGLF’s quasi-linear saturation rule is tuned to the results of a finite set of gyrokinetic runs means that the parameter space in which TGLF is valid is quite constrained. If one wishes to simulate a plasma condition that is far away from where TGLF was tuned, one should almost certainly use a fully nonlinear gyrokinetic simulation. One could then re-tune TGLF’s saturation rule with this new data, and subsequently run TGLF in this parameter space.

While the description of turbulent transport models in this section is far from complete, it should provide the reader with enough background to understand the motivation and results in the remainder of this thesis. There are, of course, far simpler models of turbulent transport in tokamaks (such as the model used in the Tokamak Simulation Code, described in Appendix F), but the gyrokinetic and gyrofluid models described here are the main focus of this thesis. This chapter now moves on to the experimental facilities used in this thesis, and then to how experimental measurements are used to validate the results of turbulent transport simulations.

2.3 Tokamaks and Standard Diagnostics in this Thesis

Experimentally, this thesis utilizes data from two tokamaks, Alcator C-Mod and ASDEX Upgrade, which are both briefly described in this section. Both machines are equipped with extensive arrays of diagnostics that measure many different properties of the plasma. In addition to specific turbulence diagnostics and data analysis techniques that were developed and used as part of this thesis, which are given considerably greater attention both later in this chapter and in Chapters 3 and 4, this thesis also made use of data from a number of more standard diagnostics on both machines. These diagnostics measure equilibrium parameters that are relevant to both turbulence measurements and to turbulent transport modeling. Specifically, turbulent transport modeling requires measurements of the magnetic equilibrium, electron and ion temperatures, electron density, plasma rotation, effective plasma charge, and radiated power. In addition, as was described in the last section, the gradients of many of these quantities are of particular importance. Tables of these inputs are shown for each of the gyrokinetic simulations in Chapter 5. These diagnostics are briefly described in this section, along with references for the more interested reader.

2.3.1 Alcator C-Mod

The first tokamak studied in this thesis was Alcator C-Mod, which is a compact, high-field, diverted machine located at the Massachusetts Institute of Technology in Cambridge, Massachusetts, USA [31]. The machine has a major radius of 0.67 m and a minor radius of approximately 0.22 m. Alcator C-Mod can run with magnetic fields between 2.7 and 8.0 T, though the standard field is 5.4T, and plasma currents between 0.25 and 2.0 MA. Up to 6 MW of ion cyclotron heating and approximately 1 MW of lower hybrid current drive are available in addition to ohmic heating. The machine has molybdenum and tungsten plasma facing components, and is periodically coated with boron for higher performance operation.

The electron temperature profile on Alcator C-Mod is measured with both Thomson Scattering [32] and an electron cyclotron emission (ECE) grating polychrometer (GPC) [33, 34]. Together these give an uncertainty in the electron temperature of approximately 10%, and, after profile fitting, an uncertainty of roughly 22% in the normalized electron temperature gradient scale length in the core plasma.

The electron density is measured primarily with the same Thomson Scattering diagnostic used to measure the electron temperature, and is complemented by line-averaged measurements with a two-color interferometer [35]. The uncertainty in the density profile is approximately 15%, and the normalized density gradient scale length uncertainty is approximately 30%.

Both the ion temperature profile and the plasma rotation profile are measured with an X-ray imaging crystal spectrometer (XICS, which is also known as HiReX on Alcator C-Mod) [36]. In general, ion temperature and rotation measurements are considerably more difficult than electron measurements, and so the ion temperature uncertainty is approximately 15%, and the rotation uncertainty can be up to 100%, depending on the magnitude of the rotation in a given discharge. The normalized

ion temperature gradient scale length uncertainty is approximately 40% after profile fitting.

The effective charge Z_{eff} is measured with a visible Bremsstrahlung spectrometer, with an uncertainty of approximately 30% [37]. The radiated power profile is measured with a bolometer array, with an uncertainty of approximately 15% [38]. Finally, the magnetic equilibrium is reconstructed with the EFIT code [39].

2.3.2 ASDEX Upgrade

ASDEX Upgrade is a medium size, moderate-field, diverted tokamak located at the Max Planck Institute for Plasma Physics in Garching, Germany [40]. The tokamak has a major radius of 1.65 m and a minor radius of approximately 0.5 m. The maximum toroidal magnetic field is 3.1 T, though the machine generally operates at around 2.5 T. Plasma currents range from 0.4 to 1.6 MA. Plasma heating consists of up to 20 MW of neutral beam injection, 6 MW of ion cyclotron heating, and 4 MW of electron cyclotron heating. Most plasma facing components are made of tungsten, and are also occasionally coated with boron.

Like on Alcator C-Mod, the electron temperature profile is measured with both a Thomson Scattering system [41, 42] and an ECE radiometer (radiometers and grating polychrometers measure the same ECE radiation, but with different detector systems) [43]. Combined these systems give an electron temperature uncertainty of approximately 9% and normalized electron temperature gradient scale length uncertainties of approximately 20%.

Also like on Alcator C-Mod, the electron density is primarily measured with Thomson Scattering, complemented by an interferometer [44]. The measured electron density has an uncertainty of approximately 12%, and the resulting fitted profiles have a normalized density gradient scale length uncertainty of approximately 30%.

Unlike Alcator C-Mod, ASDEX Upgrade measures both the ion temperature and

plasma rotation with a charge exchange recombination spectroscopy (CXRS) system [45]. Since this diagnostic relies on the presence of neutral particles in the plasma, it requires the use of neutral beams in order to make measurements. Fortunately, very short beam blips are usually sufficient for these measurements, enabling measurement in plasmas that are not dominantly beam heated. The downside of beam blip measurements is that the small amount of data collected generally leads to larger uncertainty than in cases with steady beams. The ion temperature uncertainty measured with beam blips is approximately 18% and the uncertainty in plasma rotation is approximately 48%. The normalized ion temperature gradient scale length uncertainty is approximately 30%.

Like on Alcator C-Mod, the effective charge is measured with Bremsstrahlung spectrometers with an uncertainty of 30% [46], and the radiated power profile is measured with a series of bolometer arrays, with an uncertainty of 50% [47]. The magnetic equilibrium is reconstructed with the CLISTE code [48], which also has its own uncertainties.

2.4 Validation and Validation Constraints

Armed now with a general understanding of turbulence in fusion plasmas, a few specific turbulent transport models, and two tokamaks for experimental data, it is possible to discuss the process of validation and the experimental measurements that can be used as constraints.

As has already been discussed in this thesis, one of the major motivations for developing turbulent transport simulations is to predict the performance of future fusion devices. Before one uses a turbulent transport model to predict a future device, however, one must be confident that the model accurately reproduces what is happening in current machines. The process of checking whether a turbulent transport model is correctly reproducing experimental results is that of validation. This contrasts with

what is commonly called ‘verification,’ which involves checking that a particular code has accurately implemented a given model [49]. Even if one is interested in turbulent transport models only for their insight into fundamental turbulence physics, one must use a validated model, or the insights provided do not represent what is actually happening in real plasma turbulence. The related process of choosing the correct model for a particular task is known as ‘model selection.’

In this thesis, validation is defined as [1]:

The process of determining the degree to which a model is an accurate representation of the real world from the perspective of the intended uses of the model.

This section describes validation of turbulent transport models in general, and then describes the four validation constraints that are applied in this thesis. More details about best practices in validation of turbulent transport models can be found in Reference [49].

To begin, consider the definition of validation from above: one would like to compare the results of a given model to “the real world.” In the case of experimental fusion research, and really in any endeavor, “the real world” is perhaps too strong of a statement, as the best that one can do is to make measurements of some system and quantify how closely these measurements represent “the real world” (in other words, quantify the uncertainty of the measurements). To that end, one would like to make many experimental measurements of fusion plasmas that one can then compare to the outputs of turbulent transport models, in order to validate these models.

Consider now the statement, “from the perspective of the intended uses of the model.” If one is interested in the fundamental physics of turbulence in fusion plasmas, then the intended uses of the model are manifold, and one is interested in basically any measurement that one can make of the plasma: heat fluxes (the amount of heat transported through the electron and ion components of the plasma), fluctuation

amplitudes, relative fluctuation phases, turbulent eddy sizes, etc. If, on the other hand, one is only interested in performance prediction, it may seem that one is only concerned with heat fluxes, as those are what is required in order to predict the temperature profile of a future machine given the machine parameters.

It turns out however, that comparing against many experimental measurements, which will be called ‘validation constraints’ when used for the purpose of validation, is vital even if the ultimate goal is performance prediction. In some literature, validation constraints are also called ‘quantities of interest.’ The essential reason for this is that one must ensure that the model is not giving the right answer for the wrong reasons, in what is known as ‘fortuitous agreement.’ Due to the fact that validation will only be performed in some finite parameter space, it is possible that at any given parameter point the model correctly predicts the electron and ion heat fluxes, though it does so without truly capturing the fundamental nature of the turbulence. If one were then to move to a different regime, such as the plasma in a future machine, it is possible that the model would no longer give the correct answer, as the model is no longer matching the heat fluxes by chance.

One can reduce the chances of this fortuitous agreement in a number of ways. First, one can simply validate the code in many possible conditions, increasing the parameter range over which the model would have to be correct by chance. The downside of this approach is that it requires many, many validation studies, which can be time consuming, and still leaves a finite parameter space over which the model is validated. On the other hand, one can also compare many experimental measurements to the outputs of each simulation. In this manner, one checks more fundamentally whether or not the model is correctly reproducing the underlying turbulent phenomena. As one adds more validation constraints, one increases the likelihood that the model is giving the right answer for the right reasons, and therefore the confidence that the model will still be correct when expanding to a new parameter regime.

In particular, this thesis is concerned with four validation constraints, two of which were developed as part of this thesis work. The first two are the electron and ion heat fluxes, which are the absolute minimum that one must check in order to validate any transport model. The second two, to which this thesis work made novel contributions, are the perturbative thermal diffusivity and electron temperature fluctuations. A few other validation constraints are also mentioned throughout this thesis (such as density fluctuation measurements, made in Appendix E; radial correlation lengths of electron temperature fluctuations, and the phase angle between temperature and density fluctuations, to which this thesis also made contributions and are discussed in Chapter 5), but these four are the main focus, and are now described in more detail.

2.4.1 Power Balance Heat Fluxes

The first validation constraint described in this thesis, and perhaps the most commonly applied, is the heat flux flowing through the plasma. As described above, turbulence leads to transport of both heat and particles. It is precisely this transport that makes achieving net energy production in a fusion device so difficult. While both heat and particle fluxes are important to the final plasma performance, this thesis focuses primarily on heat fluxes, as particle fluxes are a topic of considerable depth in their own right. The interested reader is directed to Reference [50] for an in depth discussion of particle fluxes in tokamaks and their measurements.

Returning to the heat flux, energy may flow through both the electron and ion species within the plasma. For this reason, the total heat flux is generally divided into the electron heat flux and the ion heat flux. Heat flux is typically quantified as either a power per unit area, which in this thesis is denoted as Q_s for species s , with units of MW/m², or as a total power through a given flux surface, which in this thesis is denoted as P_s , with units of MW (one multiplies the heat flux per unit area by the surface area of the flux surface to get the total power). In general, this heat flux may

consist of both conductive and convective components, though often the conducted power is much larger than the convected power.

In practice, however, it is impractical to directly measure the heat flux in all but the lowest performance plasmas, as no physical probe would survive the conditions inside of a fusion plasma, and there is no other good way to directly measure the heat flux. Instead, one calculates the heat fluxes using a procedure known as ‘power balance.’ Essentially, one balances known sources and sinks of energy at a given location in the plasma, and, assuming steady state, any remaining power must be transported via fluid heat flux. For this reason, one occasionally says that the heat fluxes in tokamak plasmas are ‘inferred,’ rather than measured.

More quantitatively, if one takes the energy fluid moment of Equation 2.1 (the Boltzmann Equation), one finds [51]:

$$\frac{3}{2}n_s \frac{dT_s}{dt} = C_s - \overline{\overline{P}}_s : \nabla \vec{u}_s - \nabla \cdot \vec{q}_s + S_{E,s} \quad (2.6)$$

where n_s is the density, \vec{u}_s is the fluid velocity, $\overline{\overline{P}}_s$ is the pressure tensor, C_s is the energy exchange due to collisions with other species, \vec{q}_s is the local conductive heat flux, and $S_{E,s}$ is the sum of heat sources and sinks.

Equation 2.6 says that a change in stored energy is balanced by collisional energy exchange, compression (and viscosity), conduction, and heat sources and sinks. Sources of energy include ohmic heating, ion cyclotron resonance heating (ICRH), electron cyclotron heating (ECH), neutral beam injection (NBI), and possibly alpha heating in a true fusion system. Sinks of energy include radiated power (primarily via bremsstrahlung). These terms will now all be rewritten in order to simplify the power balance process. In this section, S refers to a volumetric heating power density (MW/m³), and as before P refers to a total power and Q to a power flux density.

Specifically, define the stored energy density $W_s = (3/2)n_s T_s$ and the convected power $S_{conv,s} = \vec{u}_s \cdot \nabla W_s$, and rewrite the first term in Equation 2.6 as (assuming

density is approximately constant):

$$\begin{aligned}
\frac{3}{2}n_s \frac{dT_s}{dt} &= \frac{3}{2}n_s \left(\frac{\partial T_s}{\partial t} + \vec{u}_s \cdot \nabla T_s \right) \\
&= \frac{\partial}{\partial t} \left(\frac{3}{2}n_s T_s \right) + \vec{u}_s \cdot \nabla \left(\frac{3}{2}n_s T_s \right) \\
&= \frac{\partial W_s}{\partial t} + S_{conv,s}
\end{aligned} \tag{2.7}$$

Then, define the compressive heating as $S_{comp,s} = -p_s \nabla \cdot \vec{u}_s$, where p_s is the scalar pressure. Neglecting viscosity:

$$\vec{P}_s : \nabla \vec{u}_s = p_s \nabla \cdot \vec{u}_s = -S_{comp,s} \tag{2.8}$$

Define also $S_{cond,s} = \nabla \cdot \vec{q}_s$ and S_{ci} as the collisional heat transfer from electrons to ions.

Finally, plug in the specific heating sources and sinks [51]:

$$S_{E,e} = S_{RF,e} + S_{NBI,e} + S_{OH} - S_{rad} - S_{ion} \tag{2.9}$$

and

$$S_{E,i} = S_{RF,i} + S_{NBI,i} - S_{cx} + S_{i,s} \tag{2.10}$$

where, S_{RF} (radio frequency) is either ICRH or ECH power, S_{NBI} is NBI heating power, S_{OH} is ohmic heating power, S_{rad} is radiated power, S_{ion} is power lost to the ionization of neutrals, S_{cx} is power lost to charge exchange reactions, and $S_{i,s}$ is net power gain from the ion particle source.

Collecting these terms and rearranging to solve for the conducted power, one can

write:

$$S_{cond,e} = -\frac{\partial W_e}{\partial t} - S_{conv,e} + S_{comp,e} - S_{ei} + S_{RF,e} + S_{NBI,e} + S_{OH} - S_{rad} - S_{ion} \quad (2.11)$$

and

$$S_{cond,i} = -\frac{\partial W_i}{\partial t} - S_{conv,i} + S_{comp,i} + S_{ei} + S_{RF,i} + S_{NBI,i} - S_{cx} + S_{i,s} \quad (2.12)$$

Volume integrating these gives:

$$P_{cond,e} = -\frac{\partial U_e}{\partial t} - P_{conv,e} + P_{comp,e} - P_{ei} + P_{RF,e} + P_{NBI,e} + P_{OH} - P_{rad} - P_{ion} \quad (2.13)$$

and

$$P_{cond,i} = -\frac{\partial U_i}{\partial t} - P_{conv,i} + P_{comp,i} + P_{ei} + P_{RF,i} + P_{NBI,i} - P_{cx} + P_{i,s} \quad (2.14)$$

where $U_s = \int W_s dV$ and $P_s = \int S_s dV$.

Assuming steady state ($\partial/\partial t = 0$), collecting the conducted and convected powers into a single term $P_s = P_{cond,s} + P_{conv,s}$, dividing the powers by the area of the flux surface of interest, and neglecting terms that are typically small in the core of a tokamak ($Q_{comp,s}$, Q_{ion} , Q_{cx} , and $Q_{i,s}$)² gives the final equations for the heat fluxes:

²These terms are neglected here for the purpose of brevity, but real calculations generally do take these terms, as well as time dependence, into account.

$$Q_e = Q_{RF,e} + Q_{NBI,e} + Q_{OH} - Q_{ei} - Q_{rad} \quad (2.15)$$

and

$$Q_i = Q_{ei} + Q_{RF,i} + Q_{NBI,i} \quad (2.16)$$

In order to calculate the ion and electron heat fluxes in a tokamak plasma, therefore, one must measure or calculate the terms on the right hand sides of Equations 2.15 and 2.16, and then solve these equations throughout the plasma radius.

While it is in theory possible to do all of this analytically, the complicated physics of neutral beam, ICRH, and ECH power deposition, collisional energy transfer, and especially the plasma geometry, mean that these calculations are almost always done numerically. There are several tools that can perform these calculations, one of the most common of which is TRANSP [52]. TRANSP is used extensively throughout this thesis, for both calculating heat fluxes and for preparing inputs to turbulent transport models.

In order to calculate the terms on the right hand side of Equations 2.15 and 2.16, TRANSP requires many of the same inputs as turbulent transport models. For example, the ion and electron temperature profiles are required to calculate every term on the right hand sides. This process also requires the density profile, the magnetic geometry, the radiated power profile, and the effective charge Z_{eff} . The diagnostics used to measure all of these quantities were all described in Section 2.3.

In addition to the resulting values of the heat fluxes, it is also important to quantify the uncertainty in these values, as knowing the uncertainty is key to the validation process. One can calculate the uncertainty in the power balance outputs by propagating the uncertainty in the inputs through the governing equations. This process

gives:

$$\frac{\sigma_{Q_e}}{Q_e} \approx \left[\left(\frac{Q_{OH}}{Q_e} \right)^2 \left(\frac{\sigma_{Q_{oh}}}{Q_{OH}} \right)^2 + \left(\frac{Q_{RF}}{Q_e} \right)^2 \left(\frac{\sigma_{Q_{rf}}}{Q_{RF}} \right)^2 + \left(\frac{Q_{NBI}}{Q_e} \right)^2 \left(\frac{\sigma_{Q_{nbi}}}{Q_{NBI}} \right)^2 + \left(\frac{Q_{rad}}{Q_e} \right)^2 \left(\frac{\sigma_{Q_{rad}}}{Q_{rad}} \right)^2 + \left(\frac{Q_{ei}}{Q_e} \right)^2 \left(\frac{\sigma_{Q_{ei}}}{Q_{ei}} \right)^2 \right]^{1/2} \quad (2.17)$$

and

$$\frac{\sigma_{Q_i}}{Q_i} \approx \left[\left(\frac{Q_{RF}}{Q_i} \right)^2 \left(\frac{\sigma_{Q_{rf}}}{Q_{RF}} \right)^2 + \left(\frac{Q_{NBI}}{Q_i} \right)^2 \left(\frac{\sigma_{Q_{nbi}}}{Q_{NBI}} \right)^2 + \left(\frac{Q_{ei}}{Q_i} \right)^2 \left(\frac{\sigma_{Q_{ei}}}{Q_{ei}} \right)^2 \right]^{1/2} \quad (2.18)$$

where all variables are defined as above, σ_x is the uncertainty in variable x , and the approximate sign represents that these formulas only contain the dominant terms of uncertainty (neglecting, for example, compressional heating as described above).

Therefore, in order to calculate the uncertainty in the heat fluxes, one must calculate the uncertainties for each of the other terms in both equations. The radiated power is measured directly, so this comes directly from diagnostic uncertainty. Radio frequency power deposition uncertainty depends on exactly what type of wave is being used to heat the plasma, with ICRH having considerably larger uncertainty than ECH. The ICRH deposition power on Alcator C-Mod has an uncertainty of approximately 10%, though modeling this deposition uncertainty can be challenging. Outside of its narrow deposition region (which is the case for all of the locations in discharges analyzed in this study), the ECH on ASDEX Upgrade deposits nearly zero power, with no significant uncertainty. NBI power deposition depends on the profiles and will be calculated by the power balance solver. Note that in this thesis, none of the plasmas that are analyzed use neutral beam heating, and so this term will be zero in all cases in this thesis. The ohmic power and the electron-ion collisional exchange terms also depend heavily on the plasma profiles, and are calculated within TRANSP.

In particular, it can be difficult to estimate the uncertainty in the electron-ion exchange term in higher density plasmas where the electron and ion temperatures are nearly equal. To calculate the uncertainty on these terms, the formulas for the ohmic and collisional exchange powers are used. The calculated uncertainty in these terms is then:

$$\left(\frac{\sigma_{Qoh}}{Q_{OH}}\right)^2 = \frac{9}{4} \left(\frac{\sigma_{Te}}{T_e}\right)^2 + \left(\frac{\sigma_{Zeff}}{Z_{eff}}\right)^2 \quad (2.19)$$

and

$$\left(\frac{\sigma_{Qei}}{Q_{ei}}\right)^2 \approx 4 \left(\frac{\sigma_{ne}}{n_e}\right)^2 + \frac{1}{4} \left(\frac{\sigma_{Te}}{T_e}\right)^2 + \left(\left(\frac{\sigma_{Ti}}{T_i}\right)^2 + \left(\frac{\sigma_{Te}}{T_e}\right)^2\right) \quad (2.20)$$

These formulas explain further why the profile and effective charge uncertainties given in Section 2.3 are so important to determining the uncertainty in the overall heat flux.

In addition, Equations 2.17 and 2.18 also reveal that the exact uncertainty for the heat flux in a given plasma depends on the mix of heating and the other heat transport channels in a given discharge, so that in reality these uncertainties vary slightly for every discharge. For example, if RF heating makes up a larger portion of the heat flux for a particular plasma, the Q_{RF}/Q_i term will be larger, and the full uncertainty calculation may give a different value.

This being said, if one considers the mid- to outer-core ($\rho_{tor} \approx 0.3 - 0.8$) of a typical moderate power, ICRH heated, L-mode discharge on Alcator C-Mod, typical values for the uncertainties are approximately 25% in Q_e and 25% in Q_i . Similarly, for a moderate power, ECH-heated, L-mode on ASDEX Upgrade, typical uncertainties are approximately 18% in Q_e and 26% in Q_i . While the uncertainty may change depending on the discharge in question, these approximate values are a good estimate for a fairly wide range of plasmas.

The electron and ion heat fluxes inferred from power balance in this manner have been used as a validation constraint for turbulent transport models for many years. Nearly every validation study of a turbulent transport model, no matter what other validation constraints it includes, will almost certainly include the heat fluxes, as these are often seen as the most important constraint if the ultimate goal is the prediction of future plasma performance. One of the major criticisms of using heat fluxes as a validation constraint, however, is that they are not measured directly, and are instead inferred from calculations based on many other measurements. It is for this reason that some hierarchies, which evaluate how directly validation constraints are measured, put heat fluxes in a place of less importance than, for example, fluctuation measurements [49]. This argument, in addition to the desire to compare many constraints and avoid fortuitous agreement, motivates the development of other validation constraints, such as those described in the remainder of this section.

2.4.2 Electron Perturbative Thermal Diffusivity

The third validation constraint used in this thesis, in addition to electron and ion heat fluxes, is the electron perturbative thermal diffusivity (referred to for most of this thesis as just the ‘perturbative diffusivity’). Rather than measuring an equilibrium property of the plasma turbulence, the perturbative diffusivity measures the manner in which the electron heat flux responds to changes in the electron temperature gradient. While using a diffusivity to represent the complex process of turbulent heat transport is clearly a significant simplification, it does offer a straightforward method of characterization and can be a good way of defining useful quantities for comparison.

Before defining the perturbative diffusivity, consider first the equilibrium state of the plasma. In a traditional diffusive model for a fluid, the heat flux is linearly proportional to the temperature gradient, and this proportionality constant is known as

the diffusivity. While this is clearly a gross oversimplification for nonlinear turbulence in a plasma, it may still be a useful way of characterizing the system. If one chooses to characterize the system in this way, one calculates what is known as the power balance diffusivity, χ_e^{PB} [53]. Quantitatively, one defines the power balance diffusivity as:

$$\chi_e^{PB} = -\frac{Q_e}{n_e \nabla T_e} \quad (2.21)$$

When one plugs in these quantities using appropriate units, the power balance diffusivity ends up having units of m^2/s , which is what one would expect for a diffusivity. Note that the negative sign is included in the equation since the temperature gradient is also negative, as the temperature decreases with increasing radius. Since the temperature gradient is essentially always negative, the negative sign is occasionally omitted.

One can also define a similar quantity for ions, χ_i^{PB} , where one replaces all of the relevant quantities with that for ions. Note that later in this document any quantities without a species subscript will refer to the *electron* channel.

In a turbulent plasma, however, the heat flux is certainly not linearly proportional to the temperature gradient, as the turbulence amplitude and heat flux are inherently nonlinear. It is for this reason that one generally uses the heat fluxes, as described in the last section, and not the power balance diffusivities, as validation constraints. On the other hand, it is useful to define a quantity that gives information about the incremental change in heat flux due to an incremental change in the temperature gradient, around an operating point. This *perturbative* electron thermal diffusivity is defined as [54]:

$$\chi_e^{pert} = -\frac{1}{n_e} \frac{\partial Q_e}{\partial \nabla T_e} \quad (2.22)$$

where all relevant quantities are defined as above. In past literature, this quantity has also been called the *heat pulse* diffusivity (since it governs the propagation of heat pulses) [54] and the *incremental* diffusivity (since it is the change in heat flux in response to an incremental change in the gradient) [53].




Figure 2-2: An illustration of the definitions of the power balance and perturbative thermal diffusivities for a plasma operating at the black point, with an arbitrary flux-gradient relationship represented by the red line. The slope of the orange line is the power balance thermal diffusivity. The slope of the green line is the perturbative thermal diffusivity. Figure originally from [55], and based on a figure from [53].

Figure 2-2 illustrates graphically the differences between χ_e^{pert} and χ_e^{PB} . This figure shows what is known as a ‘flux-gradient’ relationship, which, as the name suggests, illustrates the relationship between the heat flux and the temperature gradient for a given plasma state. In this particular figure, the plasma is assumed to be operating at the black point. The red curve represents an arbitrary flux-gradient relationship (which is determined by the complex nonlinearities of the turbulence). The slope of the orange line, which connects the origin to the operating point, is the power balance diffusivity. The slope of the green line, which is tangent to the red curve at the plasma operating point, is the perturbative diffusivity. These diffusivities are only the same if the heat flux and temperature gradient are linearly related with

no offset [53]. This is not theorized to be the case for turbulent plasma systems, which are often characterized by an offset critical gradient model (the heat flux is near zero below some critical gradient, and then roughly linearly increasing with gradient above some critical gradient, which is the gradient at which some turbulent mode is driven unstable) [4]. Note that a real flux-gradient relationship would reflect this critical gradient model, and that the curve in Figure 2-2 is meant only as an illustrative example. Appendix A describes the historical context of the difference in these definitions in further detail.

Given this definition, one can interpret the perturbative diffusivity as the degree to which the turbulence acts to resist a change in the gradient. A high perturbative diffusivity indicates that increasing the gradient even slightly increases the heat flux significantly, meaning that one would have to apply considerable additional heating in order to increase the plasma temperature even slightly. A low perturbative diffusivity indicates the opposite. It is for this reason that the perturbative diffusivity is sometimes referred to as the ‘stiffness’ of the temperature profile, as it characterizes its resistance to change [56]. References [53, 56, 57, 58, 59, 60, 61, 62, 63], as well as those in Appendix A further describe the importance of the perturbative diffusivity.

Given this understanding of what the perturbative diffusivity is and why it is important, consider now how one goes about measuring the perturbative diffusivity both experimentally and in the outputs of simulations, so that one can use it as a validation constraint. One of the main contributions of this thesis was the development of a new method with which to measure the perturbative diffusivity in tokamak plasmas, which is described in great detail in Chapter 3. The essential idea behind these measurements is that one can calculate the perturbative diffusivity of a plasma by tracking a transient heat pulse as it propagates through the plasma. Versions of this method had been used in work in the past [53, 56, 57, 58, 59, 60, 61, 62, 63], but this thesis developed a new method of making this measurement, based on tracking

heat pulses generated by partial sawtooth crashes, which is often more readily available and robust than past methods. Chapter 3 covers this topic in considerably more detail, including details of the calculation and its uncertainty.

Finally, unlike the electron and ion heat fluxes, the perturbative diffusivity is not a standard output of most turbulent transport simulations. Generally, these codes do not attempt to simulate a transient heat pulse, since this time dependent work would take an enormous amount of computational power. Instead, simulations generally try to map out a small portion of the flux-gradient relationship, shown in Figure 2-2. The electron temperature gradient that is input to the code is scanned slightly above and below the operating point in order to show how the heat flux changes in response to small changes in the temperature gradient. One then takes the slope of the line mapped out by these simulations in order to calculate the perturbative diffusivity. Since generally turbulent transport models take the normalized gradient length scales as inputs, instead of the temperature gradients directly, one can write [55]:

$$\chi_{Sim.}^{pert} = \frac{a}{n_e \cdot T_e} \frac{\partial Q_e}{\partial(a/L_{Te})} \quad (2.23)$$

where all variables are defined as above and $\chi_{Sim.}^{pert}$ is the perturbative diffusivity estimated by the simulation model.

This process is described in considerably more detail in Chapter 5, and is illustrated there in Figure 5-6. The same general process is used for all of the turbulent transport models analyzed in this thesis, though the exact implementation differs slightly depending on the experimental discharge (for example, the amount that the electron temperature gradient is scanned up and down).

Once one has both an experimental measurement of the perturbative diffusivity and a way to calculate it based on the outputs of a turbulent transport simulation, one can use the perturbative diffusivity as a validation constraint, in the same manner as

one can compare the experimental and predicted heat fluxes. This was done previously in Reference [63], as well as in References [55, 64, 65], which were published as part of this thesis work.

2.4.3 Electron Temperature Fluctuations

The final validation constraint used extensively in this thesis is the electron temperature fluctuation amplitude. As described in Section 2.1, plasma turbulence leads to fluctuations in the plasma potential, density and temperature. The different modes of plasma turbulence described in that section lead to different spatial scales and frequency spectra of fluctuations, and thus by comparing observations of these fluctuations to the fluctuations predicted by simulations, one gains insight into whether or not the code is predicting the right type of turbulence. Since temperature fluctuations are a fairly direct consequence of turbulence, they are often treated as a high fidelity constraint of turbulent transport codes.

Unfortunately, however, fluctuations are often quite difficult to measure experimentally, which is one of the reasons why they have not been used extensively as validation constraints. As with heat fluxes, it is possible to measure fluctuations with physical probes, but the extreme conditions inside of a fusion plasma eliminate this possibility. Instead, the temperature fluctuations in this thesis are measured based on gathering and analyzing electron cyclotron emission (ECE) radiation. This section gives a brief description of the relevant physics, and the interested reader should see Reference [66] for a much more thorough description of the phenomena.

Very simplistically, ECE arises from the acceleration experienced by electrons in a background magnetic field [66]. As the electron undergoes its gyro-motion, it is constantly accelerated by the Lorentz force. Since the electron is an accelerating charged particle, it will emit electromagnetic radiation just like any other accelerating

charged particle, according to the (unfortunately rather complicated) formula:

$$\frac{d^2P}{d\Omega_s d\nu} = \frac{q^2\omega^2}{4\pi\epsilon_0 2\pi c} \left| \int_{-\infty}^{\infty} \hat{R} \times \left(\hat{R} \times \frac{\vec{v}}{c} \right) \exp \left(i\omega \left(t' - \frac{\hat{R} \cdot \vec{r}}{c} \right) \right) dt' \right|^2 \quad (2.24)$$

where P is radiated power, Ω_s is solid angle, ν is outgoing radiation frequency (in Hz), q is the charge of the particle, ω is the frequency of the particle's acceleration (in radians per second), c is the speed of light, \hat{R} is the unit vector from the charge to the field point, \vec{v} is velocity of the particle, t' is the retarded time, and \vec{r} is the position of the charge. Again, the reader is directed to Reference [66] for the detailed derivation of this formula.

Considering specifically gyro-motion, which is the origin of ECE, the acceleration occurs at a very specific frequency, known as the cyclotron frequency:

$$\omega_{ce} = \frac{eB}{m_e} \quad (2.25)$$

where e is the fundamental charge, B is the background magnetic field, and m_e is the electron mass.

Plugging this (including a relativistic mass correction) into the formula above, and performing the necessary integrals over the particle orbit, one finds that the radiation emitted from the gyro-motion occurs in harmonics, and is affected by the Doppler shift of the electron as it moves in its orbit:

$$\omega = m\omega_{ce} \frac{\sqrt{1 - \beta_r^2}}{1 - \beta_{r,\parallel} \cos \theta} \quad (2.26)$$

where m is an integer representing the harmonic and $\beta_r = v/c$. The numerator of the last term is the relativistic mass correction and the denominator is the Doppler shift.

Since the dominant toroidal magnetic field in a tokamak varies as $B \sim 1/R$ due to

the toroidal geometry of the machine, where R is the major radius of the device, each radial location emits ECE at a different frequency. The Doppler shift and relativistic corrections slightly broaden the emission layer, which can be taken into account in detailed ECE calculations, though this is generally a small effect.

In this way, by choosing a frequency to observe, one can select a radial location in the plasma from which to collect ECE radiation.

Going further, if one assumes a Maxwellian velocity distribution in Equation 2.24, and then assumes that the plasma is optically thick and therefore acts as a blackbody for ECE at the location of the ECE resonance (the plasma absorbs all incident radiation at this frequency, while also emitting radiation at this frequency; in other words, the photon mean free path is short), then one finds that the intensity of the ECE at a frequency ν is:

$$I(\nu) \approx \frac{\nu^2 T_e}{c^2} \quad (2.27)$$

where T_e is the temperature of the plasma. The emitted intensity of the light is therefore directly proportional to the temperature of the plasma.

In order to behave like a blackbody and absorb incident radiation at the cyclotron frequency, a plasma must have an optical depth $\tau \gg 1$, where the optical depth goes roughly as [66]:

$$\tau \sim \int n_e T_e^{m-1} ds \quad (2.28)$$

where n_e is the plasma density, m is the ECE harmonic, and s is the path length of the radiation. Even though strictly one requires $\tau \gg 1$, in practice $\tau > 2$ is used as a cutoff.

More precisely, the optical depth has further dependencies on the ECE harmonic, as well as the polarization of the emitted light. In terms of polarization, light emit-

ted with its electric field parallel to the background magnetic field is known as the ordinary mode, or ‘O-mode.’ On the other hand, light emitted with its electric field perpendicular to the background magnetic field is known as the extraordinary mode, or ‘X-mode.’ These two modes have different optical depths and also different cutoffs and resonances within the plasma.

Most practical ECE diagnostics use the second harmonic X-mode emission, since first harmonic X-mode is always cutoff on the low-field side, first harmonic O-mode may be cutoff depending on the density, and higher harmonics of both tend to be optically thin.

For the second harmonic X-mode, the optical depth is given by [67]:

$$\tau_2^X = \pi q \frac{T_e}{m_{e0} c^2} \frac{\omega_{ce}}{c} \left(\frac{12 - 8q + q^2}{12 - 4q} \right)^{1/2} \left(\frac{6 - q}{6 - 2q} \right)^2 L_B \quad (2.29)$$

where $q = (\omega_{pe}/\omega_{ce})^2$, $L_B = B_0 |dB_0/ds|$, and all other variables are defined as above. Appendix C derives this equation from the starting point given in Reference [68].

To summarize the principle by which one can use ECE diagnostics to measure the electron temperature of the plasma, knowing the toroidal magnetic field, one selects a frequency in order to specify a radial location. Light collected at this frequency is then directly proportional to the electron temperature, giving a localized temperature measurement. Details of diagnostic design vary considerably, but in general it is possible to calibrate the diagnostic such that one can obtain an output in volts, and then convert this to keV of electron temperature with some keV per volt calibration factor.

This completes the general description of ECE diagnostics, including the GPC and radiometer diagnostics mentioned above in Section 2.3. Unfortunately, however, straightforward ECE measurements are often unable to resolve turbulent fluctuations in the electron temperature. While the photon energy from ECE is fairly small,

photon statistics are not negligible when making ECE measurements. Depending on the details of the system, photon statistics generally limit ECE measurements to accuracies of approximately 1%. Turbulent temperature fluctuations, however, are often on the order of 1%, if not smaller [66].

The temperature fluctuation diagnostic described in this thesis, correlation electron cyclotron emission (CECE), overcomes this challenge by correlating signals from two closely spaced bandpass filters. By choosing frequency ranges that do not overlap, the CECE diagnostic measures different photons coming from different plasma volumes. If the frequency ranges are close enough, however, they are still measuring the same turbulent structures, and so observe the same temperature fluctuations. With this technique, it is possible to resolve fluctuation levels as small as approximately 0.1 %. This technique has been successfully employed on a number of machines worldwide [69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79].

Chapter 4 describes in detail the CECE hardware that was constructed as part of this thesis for both Alcator C-Mod and ASDEX Upgrade [80, 81], as well as contributions that were made to the manner in which one analyzes data collected by CECE diagnostics. That chapter also gives examples of experimental measurements from both machines.

Finally, consider how one extracts the electron temperature fluctuation level from the outputs of gyrokinetic or gyrofluid simulations. Regardless of how exactly the code outputs temperature fluctuation data, a point measurement contains information from all spatial scales in the simulation. The diagnostic, however, only measures fluctuations on some finite spatial scale, since fluctuations smaller than the spot size of the optical system are averaged out by the measurement volume. For this reason, one must somehow filter the simulation output in order to account for the measurement characteristics of the real diagnostic. This process is known as applying a ‘synthetic diagnostic,’ and is described in more detail in Chapters 5 and 6 for the

specific codes used in this thesis.

Conclusions

This chapter has introduced the general concepts of turbulence and turbulent transport in tokamak plasmas, a few models used to analyze this turbulence, the tokamaks studied in this thesis, the concept of validation, and the validation constraints used in this thesis. While the details of turbulence theory are not important to understanding the remainder of this thesis, the turbulent modes and some general characteristics of turbulence described in Section 2.1 will be important in particular to understanding the results of Chapters 5 and 6. The three turbulent transport simulation codes described in Section 2.2 are also key to these two chapters. The general tokamak information in Section 2.3 comes up throughout this thesis, as it is relevant to both the experimental measurements and validation efforts described here. The contributions that this thesis made to developing the perturbative diffusivity validation constraint are the focus of Chapter 3, and the contributions to the hardware and analysis of CECE data are the focus of Chapter 4. The heat flux validation constraint will come up throughout the two validation-focused chapters.

Chapter 3

Perturbative Thermal Diffusivity From Partial Sawtooth Heat Pulses

This chapter describes a method of measuring plasma perturbative thermal diffusivity based on tracking heat pulses generated by partial sawtooth crashes, which was developed as part of this thesis. In addition to purely experimental studies, the perturbative thermal diffusivity serves as a key constraint in the validation of turbulent transport models, motivating its measurement in as many plasma conditions as possible. While the perturbative diffusivity has been measured on tokamaks before, the method based on partial sawtooth heat pulses is a unique contribution of this thesis work. One significant advantage of this method as compared to the more established modulated electron cyclotron heating method is that no active external perturbation of the plasma is necessary, providing measurement of another plasma parameter without any additional hardware or changes in machine operation.

The remainder of this chapter describes what partial sawteeth are, explains how one calculates the perturbative thermal diffusivity from the propagation of partial sawtooth heat pulses, compares this new method with both an obsolete method based on full sawteeth and the more established method of measuring perturbative

diffusivity based on modulated electron cyclotron heating, and presents experimental observations from Alcator C-Mod and ASDEX Upgrade.

3.1 Partial Sawteeth

Perturbative transport studies, in which one observes how some type of perturbation affects or propagates in a tokamak plasma, have been used for many years on tokamaks in order to provide insight into fundamental plasma behaviors. References [53] and [56] give a fairly thorough background of perturbative transport studies, including studies based on both heat and cold pulses. These heat and cold pulses can be generated in a variety of ways, including local modulated electron cyclotron heating deposition and sawtooth crashes to generate radially propagating heat pulses, and intentional impurity injection for edge cooling to generate cold pulses. While References [53] and [56] lay out a broader context for these methods, work on Alcator C-Mod in particular includes References [34] for sawteeth and [82] for impurity injection. The present work will focus on heat pulses generated by partial sawtooth crashes and their utility in measuring plasma perturbative thermal diffusivity.

In order to properly calculate a perturbative thermal diffusivity from the propagation of heat pulses, the heat pulses must be moving through the plasma in a diffusive manner [55]. Specifically, one must make sure that the heat is not being transported by phenomena other than standard turbulent heat transport, such as macroscopic plasma motion caused by MHD activity. In particular, this requirement has made the use of heat pulses generated by full sawtooth crashes, which were for many years used to measure perturbative thermal diffusivity [57], rather difficult. A variety of work showed that heat pulses generated by sawtooth crashes were often accompanied by non-diffusive “ballistic” transport [58]. Figure 3-1 (a) illustrates a sawtooth crash and ballistic transport. This ballistic transport is thought to be related to the MHD phenomena that cause the sawtooth crash in the first place, and may extend quite




Figure 3-1: Example temperature profiles of a tokamak plasma before and after (a) a full sawtooth crash and (b) a partial sawtooth crash. The red line is the profile before the crash. The blue and green lines are the profile afterward. The inversion radius is the point at which the temperature is the same before and after the crash (the division between where the crash causes a temperature decrease and a temperature increase). The mixing radius is the radius out to which the crash flattens the temperature profile through fast MDH transport. (a) also illustrates what is known as the ballistic effect, where some heat is transported past the mixing radius on non-diffusive time scales.

far past the sawtooth mixing radius (the radius out to which the sawtooth crash flattens the plasma temperature profile). Another interpretation of ballistic transport is that the full sawtooth heat pulse is so large in amplitude that it changes the local temperature gradient enough that one is no longer truly measuring the local slope of the flux-gradient curve shown in Figure 2-2. Further information about the body of work leading to the conclusion that full sawtooth heat pulses are inappropriate for measuring perturbative diffusivity is given in Appendix A.

This same past work noted, however, that “partial” or “compound” sawtooth crashes generate heat pulses that are not affected by the same ballistic transport [58]. A partial sawtooth is the temperature crash of an annular region of the plasma, rather than the entire core, caused by a similar mechanism as full sawteeth. One possible explanation for the annular crash is the transient presence of multiple $q = 1$

surfaces while heat and particles (and current density) flow back into the core after a sawtooth crash. The exact physics of the partial sawtooth crash is not, however, the focus of this thesis. The significantly smaller radial extent of the crash region and smaller amplitude of the heat pulse make partial sawtooth heat pulses far more appropriate for calculating the perturbative thermal diffusivity outside of the sawtooth mixing radius. The differences between partial and full sawteeth, the manifestation of ballistic transport, and the location of the sawtooth mixing radius are all illustrated in Figure 3-1. While this figure is a cartoon representation of these phenomena, experimental measurements showing exactly these effects have been made, for example in Reference [83].

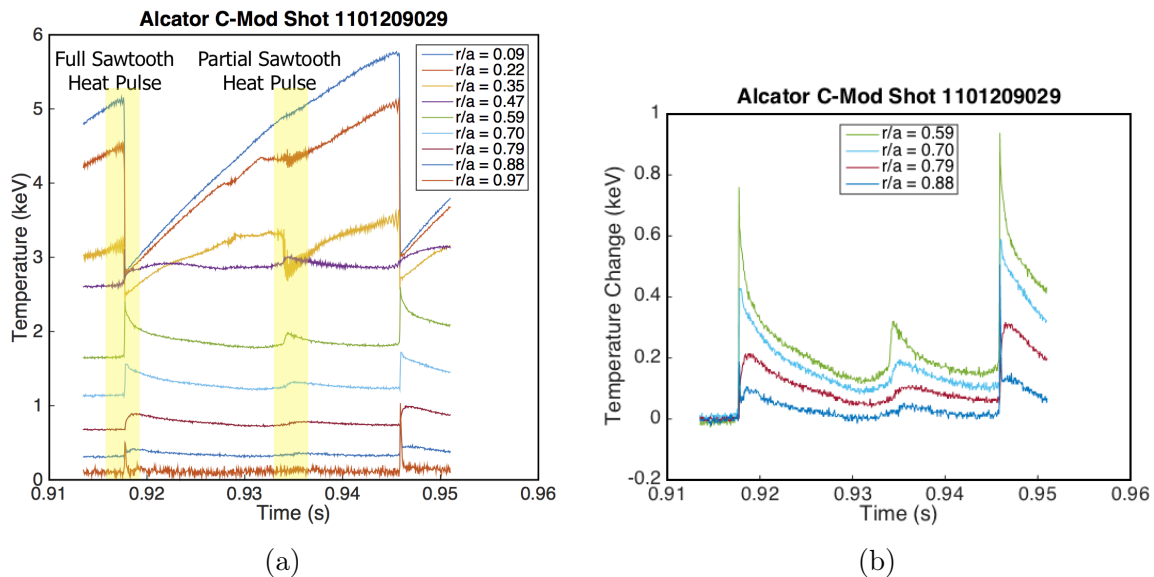


Figure 3-2: Full and partial sawteeth represented in the time domain in real Alcator C-Mod data. This data is from the nine channels of the grating polychrometer (GPC) system on Alcator C-Mod. The time trace in (a) shows both full and partial sawteeth. (b) shows only channels 5 through 8 of the GPC during the same time period, with the background temperature subtracted off for each channel. The non-diffusive full sawtooth heat pulse and diffusive partial sawtooth heat pulse are clearly evident in this data. Figure from [55].

One can also look at the full and partial sawtooth crashes in the time domain, allowing one to observe the heat pulses that they create at radii outside of the mixing radius. Figure 3-2 shows partial and full sawteeth from real data on Alcator C-

Mod, where the electron temperature was measured with a grating polychrometer (GPC) [55]. Further information on Alcator C-Mod and the GPC system was given in Chapter 2.

From a time domain perspective, the ballistic transport associated with full sawteeth manifests as a very sudden rise in the temperature at a location outside of the mixing radius. This can be seen on channels 5 and 6 (dark green and light blue) in Figure 3-2. The combination of sharp rise and gradual decay was one of the key characteristics that led past work to show that these heat pulses were inconsistent with a computational diffusive model [58]. In contrast, the partial sawtooth in Figure 3-2 shows a more gradual rise, consistent with a diffusive process. These observations indicate that partial and full sawteeth on Alcator C-Mod exhibit the same characteristics as those studied in past work [58].

3.2 Perturbative Diffusivity Calculation Method

Once a diffusive heat pulse has been generated (with partial sawteeth, as in this work, or with modulated electron cyclotron heating), the perturbative thermal diffusivity can be calculated by observing the propagation of this pulse in space and time. In particular, the velocity of the peak of the pulse and the rate at which the amplitude of the pulse damps in space are used in the calculation of the perturbative diffusivity. This work will utilize the “extended time-to-peak” method of calculating perturbative thermal diffusivity based on the propagation of heat pulses, originally presented and derived in References [54] and [84]. This method is more generally applicable than the basic time-to-peak method [84] and does not require regularly spaced pulses, as is the case with Fourier analysis [85]. Previous work has confirmed the general agreement (to within roughly 20%) between the extended time-to-peak method and the Fourier method [60, 85], so even though Fourier analysis is more common in much modern work with modulated electron cyclotron heating, the choice of method should not

impact the final results.

The central equation of the extended time-to-peak method can be written as [54]:

$$\chi_e^{pert} = 4.2a_c \frac{v_{HP}}{\alpha} \quad (3.1)$$

where:

$$v_{HP} = \sqrt{\kappa} \frac{a}{a-s} \left(\frac{dt_{peak}}{dr} \right)^{-1} \quad (3.2)$$

is the radial velocity of the peak of the heat pulse (in m/s) and:

$$\alpha = 10(a-s) \frac{d}{dr} \log(A) \quad (3.3)$$

is a parameter describing damping of the heat pulse as it propagates radially (unitless).

Other variables are defined as follows: minor radius a (m), minor radius corrected for elongation a_c (m), radius from plasma center r (m), Shafranov shift s (m), elongation κ , time that the peak of the heat pulse reaches a given radius t_{peak} (s), and heat pulse amplitude A (eV).

The derivation of this formula in Reference [54] is not machine specific, having been solved in cylindrical coordinates and then corrected for a toroidal geometry and a shaped plasma. The primary assumptions are that the heat pulse is diffusive and that it is applied outside of the sawtooth mixing radius. It has been used on various shaped tokamaks, and fully accounts for differences in machine size and shaping. The numerical prefactor is a result of the geometry of the calculation, and should not depend on the machine for which the formula is used. In order to verify that this is the case, the calculations laid out in References [54] and [84] were repeated with the exact geometry and locations relevant to the measurements made on Alcator C-Mod. These calculations consistently resulted in pre-factors of between 3.8 and 4.3, with variations

caused by the exact relationship between the mixing radius and the measurement location. This variation of less than 10% compared to the published value of 4.2 is consistent with the published uncertainty of 10%. See Appendix B for more details on these calculations. To remain consistent with past work and between calculations presented in this work, the published value of 4.2 was retained for all calculations of the perturbative thermal diffusivity. One could in principle recalculate this prefactor for the exact mixing radius and measurement location for every discharge, but this would be time consuming and would not significantly improve accuracy.

In addition, this calculation was shown to be robustly indifferent to the shape of the initial perturbation, with the prefactor varying by less than 1% with a variety of different shapes and less than 10% with experimentally relevant perturbation widths, again less than the published uncertainty of 10%. This formula should therefore be equally applicable to heat pulses generated by various different methods. Further information is given in Appendix B.

One caveat is that this method inherently calculates a radially averaged diffusivity, over the radii through which the heat pulse propagates, since turbulent transport, and thus the perturbative diffusivity, varies with radius through the plasma. This radial averaging has been suggested as a possible cause of minor disagreements with the Fourier method of calculating the perturbative thermal diffusivity, since the mathematical mechanics of the two methods will weight the radial average slightly differently [60].

The minimum distance over which the measurement must be averaged is set by the separation of temperature measurement channels. Since high time resolution is required, electron cyclotron emission (ECE) diagnostics have been used thus far for these measurements. The same ECE diagnostics on Alcator C-Mod and ASDEX Upgrade used to measure the electron temperature profile, described in Chapter 2, were also used to measure the electron temperature for the perturbative diffusivity

measurements. Since the details of these diagnostics are relevant to the limits on the perturbative diffusivity measurement, a slightly more in depth description is given here.

The grating polychrometer (GPC) on Alcator C-Mod [33, 34] collects second harmonic X-mode electron cyclotron emission from the plasma and then splits the spectrum using a diffraction grating (see Reference [66] for a first principles description of electron cyclotron emission). The system has 9 channels at different radial locations (achieved by filtering different frequencies) spaced by approximately 2 cm. This system can sample up to 100 kHz, though it more typically runs at 20 kHz, as it was for all discharges in this study.

The heterodyne ECE radiometer on ASDEX Upgrade [43] also collects second harmonic X-mode electron cyclotron emission, but uses a series of bandpass filters to split the incoming light into frequencies. The system has 24 core channels, spaced by approximately 1.2 cm. The system can sample at up to 1MHz, though it typically runs at lower sampling rates.

On both machines, the spacing of ECE channels sets the minimum perturbative diffusivity resolution to a few centimeters. Including more than just two channels does, however, increase the reliability of the measurement, as it minimizes the impact of single channel noise on the result.

The radial range over which the measurement can be made is constrained by two separate effects. The minimum radius at which one can apply this method is the mixing radius of the partial sawtooth, as shown in Figure 3-1. Anything inside of the mixing radius (or very close outside) will have transport dominated by MHD effects related to the sawtooth crash itself, and therefore is not an appropriate region in which to measure a diffusivity. The maximum radius at which the measurement can be made is typically set by the noise level of the temperature diagnostic and the amplitude of the pulse. At some point, the heat pulse has decreased enough in

amplitude that it is no longer distinguishable from noise.

The combination of these two constraints has limited most measurements on Alcator C-Mod to the range of $0.6 \leq \rho_{tor} \leq 0.9$ [55], and measurements on ASDEX Upgrade to $0.2 \leq \rho_{tor} \leq 0.6$ [64]. The difference is due to the higher current density and thus larger sawtooth mixing radius on Alcator C-Mod, which moves the minimum radius outward, but also creates larger pulses, which are more easily measurable at radial positions further out.

The data shown in Figure 3-2 clearly contains some level of diagnostic noise in all channels. This level of noise tends to be fairly small on the grating polychrometer (GPC) system on Alcator C-Mod (much smaller than the pulse amplitude), but can be nearly of the order of the pulse amplitude on the ECE radiometer on ASDEX Upgrade (see Figure 3-3 later for an example of ASDEX Upgrade data). In any case, the time that the peak of the heat pulse reaches a given radius t_{peak} used in Equation 3.2 should reflect the actual peak of the pulse amplitude, not noise on top of the pulse. To this end, the raw temperature data is smoothed in order to remove this noise.

While the exact method of smoothing should not impact the final calculation, the smoothing must sufficiently remove variations due to noise without smoothing so much as to alter the overall shape of the pulse. This proves a tricky balance, especially considering the innermost channels, which tend to have a very fast rise in temperature at the beginning of the pulse. For this reason, in the code implementation of this calculation, different smoothing levels were used for different temperature channels, with inner radial channels being smoothed less than outer channels. This method successfully retained the fast rise of the inner channels while smoothing out noise that was a larger fraction of the amplitude in outer channels. One example of smoothing of particularly noisy data from ASDEX Upgrade is shown later in Figure 3-3 (a).

In particular, two primary smoothing algorithms were used in this work. Initially

the code to calculate the perturbative thermal diffusivity was written in MATLAB.¹ This code, which was used in Reference [55], used the MATLAB ‘smooth’ function, which is based on a “local regression using weighted linear least squares and a 2nd degree polynomial model.”² A second iteration of the code, which was used for the data processing in Reference [64], was written in Python.³ This code used a ‘lowess’ smoother, which stands for ‘Locally Weighted Scatterplot Smoothing.’⁴ These two smoothing methods were directly compared on the same data set (on Alcator C-Mod) and were shown to agree within experimental uncertainty (which is discussed further below).

In addition, due to the somewhat noisy nature of any temperature measurement, the measurement of perturbative thermal diffusivity will be averaged over many partial sawtooth crash-generated heat pulses throughout some steady state portion of a discharge (anywhere between 3 and 40 pulses for the data shown in this work). There are two ways of doing this averaging. In many past studies, this averaging was accomplished by combining many heat pulses into a single composite sawtooth, synchronized by the crash time [54]. In essence, the time traces of the temperatures after the crash were averaged for each channel, to create a single set of time traces for a superimposed pulse, and then the perturbative diffusivity was calculated for this single composite pulse.

In addition, one can calculate the perturbative thermal diffusivity for each pulse individually, and then average the values obtained for each pulse. In general, this method requires slightly less noisy data, as one would like to get at least a reasonable value for each individual pulse, but comes with the advantage that one can better quantify the uncertainty in the measurement, as one calculates a set of values rather

¹See <http://www.mathworks.com/products/matlab/> for documentation concerning MATLAB and its functions.

²Ibid.

³See <https://www.python.org/> for information concerning Python.

⁴See http://www.statsmodels.org/0.6.1/generated/statsmodels.nonparametric.smoothers_lowess.lowess.html for discussion of the particular implementation of this lowess filter.

than just one. This method also allows some level of time resolution through a non-steady shot, though the scatter can be quite large.

Nominally, these two methods should agree to within experimental uncertainty, as both are using the same set of data to calculate the same value. To confirm this agreement, both calculation methods were used to calculate the perturbative thermal diffusivity for a set of discharges on ASDEX Upgrade [64].

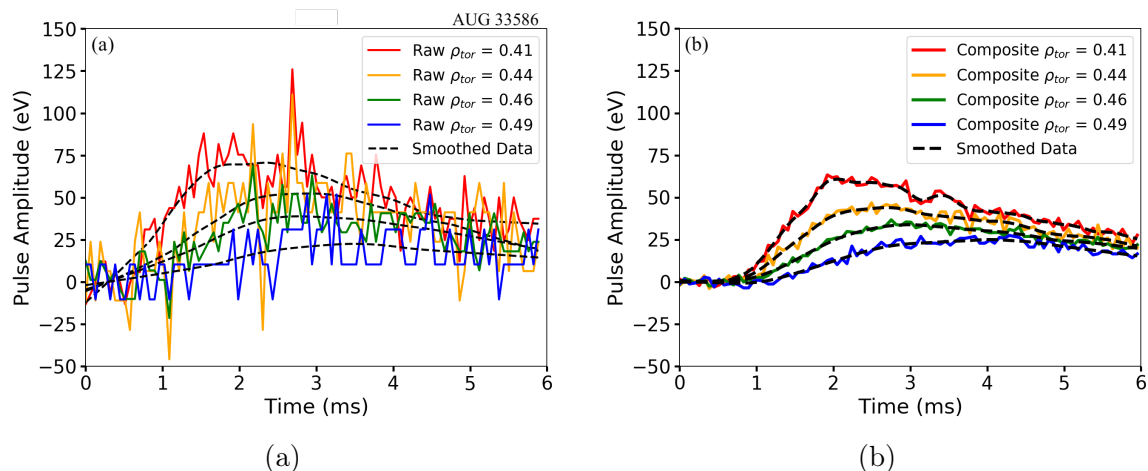


Figure 3-3: Comparison of (a) a single partial sawtooth-generated heat pulse and (b) a composite heat pulse formed by combining data from 36 individual pulses, including the pulse shown in (a). Data is taken from ASDEX Upgrade. Figure from [64].

Figure 3-3 shows a single pulse generated by a partial sawtooth crash and a composite heat pulse constructed from 36 individual pulses in one plasma discharge, including the single pulse shown. The time of peak amplitude on the innermost radial channel (the channel with the largest amplitude in Figure 3-3) was used to synchronize the heat pulses for averaging.

Qualitatively, it is immediately clear that the composite data shows the propagation of the heat pulse much more clearly than the single pulse. One may even question whether any information is salvageable from single pulses, as the data looks particularly noisy. It turns out, however, that the two methods give the same result within experimental uncertainty (see below for discussion of how this is calculated). In the discharge shown in Figure 3-3, averaging 36 individual heat pulses gives a result of

$\chi_e^{pert} = 5.04 \pm 1.35 \text{ m}^2/\text{s}$, while performing the calculation on a composite sawtooth formed by averaging the 36 pulses together gives $\chi_e^{pert} = 4.17 \text{ m}^2/\text{s}$ (the standard error between pulses cannot be calculated for the composite pulse, so only the total uncertainty for the averaging of pulses is shown). While clearly not identical, these two calculations agree to within the experimental uncertainty.

Applying both methods to a larger set of data from 24 discharges on ASDEX Upgrade reveals that the composite heat pulse method and averaging individual heat pulses agree within one standard deviation 67% of the time (for 16 of the 24 discharges), which is consistent with a normal distribution of error. Given this result, these two methods can be used interchangeably.

The experimental uncertainty has now been mentioned several times, meriting discussion of how it is calculated. There are two primary types of error that characterize the perturbative thermal diffusivity calculated by averaging many individual pulses. First, there is diagnostic error on the inputs to Equation 3.1. These include uncertainties in the temperature, as measured by whichever electron temperature diagnostic is being used in the calculation, and uncertainty in the radial location of the measurement, primarily determined by the diagnostic mapping. Since it is only the relative distance between measurement locations that is important, not the absolute radius, this error may be smaller than one would expect for a typical equilibrium mapping uncertainty. There is also some uncertainty in the exact time that the peak of the heat pulse reaches a given radius, due both to the finite time resolution of the diagnostic and due to the smoothing process described above.

These experimental uncertainties are calculated as follows:

$$\frac{\sigma_{diag}}{\chi_e^{pert}} = \sqrt{\left(\frac{\sigma_\alpha}{\alpha}\right)^2 + \left(\frac{\sigma_v}{v_{HP}}\right)^2} \quad (3.4)$$

where:

$$\frac{\sigma_v}{v_{HP}} = \sqrt{\left(\frac{\sigma_x}{\Delta x}\right)^2 + \left(\frac{\sigma_t}{\Delta t}\right)^2} \quad (3.5)$$

and

$$\frac{\sigma_\alpha}{\alpha} = \sqrt{\left(\frac{\sigma_x}{\Delta x}\right)^2 + \left(\frac{\sigma_{\log A}}{\Delta \log A}\right)^2} \quad (3.6)$$

In these equations, σ_x is the uncertainty in the difference of the radial position between two adjacent measurement channels, Δx . It is the uncertainty in the channel spacing that matters, not the absolute position, since the perturbative diffusivity calculation utilizes the separation, not the absolute position. Similarly, σ_t is the uncertainty in the time that the peak of the heat pulse reaches a given location and Δt is the time elapsed between the peak reaching subsequent channels. Finally, as stated above, A is the amplitude of the pulse, and so $\sigma_{\log A}$ is the uncertainty in the logarithm of the amplitude (which is equal to σ_A/A).

As mentioned above, on Alcator C-Mod the heat pulses were tracked with GPC [33, 34], which was described in greater detail in Chapter 2. On ASDEX Upgrade, the heat pulses were tracked with an ECE radiometer [43], which was also described in more detail in Chapter 2.

In addition to the diagnostic uncertainty, when averaging many heat pulses together one can also calculate a statistical uncertainty (standard error) for the set of measurements. This standard error is given as:

$$\sigma_{stderr} = \frac{\text{StandardDeviation}(\chi_{pulse}^{pert})}{\sqrt{N_{pulse}}} \quad (3.7)$$

where the numerator is the standard deviation of the perturbative thermal diffusivities calculated for each individual heat pulse, and N_{pulse} is the total number of pulses.

Taking these two errors together, one can calculate the total uncertainty on the measurement of the perturbative thermal diffusivity as:

$$\sigma_{\chi} = \sqrt{\sigma_{diag}^2 + \sigma_{stderr}^2} \quad (3.8)$$

On Alcator C-Mod, the diagnostic error is typically approximately 15% and the standard error varies, but for a value of 20% the total uncertainty would be $\sigma_{\chi} = 25\%$. ASDEX Upgrade typically has a comparable level of uncertainty.

This completes the description of how the perturbative thermal diffusivity of a tokamak plasma is calculated from the propagation of partial sawtooth-generated heat pulses. This method will now be compared to other methods of measuring the perturbative thermal diffusivity.

3.3 Comparison to Full Sawtooth Heat Pulses and Modulated Electron Cyclotron Heating

This section will compare measurements of the perturbative diffusivity made with partial sawteeth to measurements made with other methods. First, the use of partial sawtooth-generated heat pulses will be compared to the use of full sawtooth-generated heat pulses. This comparison was made originally in Reference [55]. As discussed above, since full sawtooth-generated heat pulses are known to exhibit ballistic transport and are therefore not good candidates for a diffusivity measurement [58], this comparison is included *purely for completeness and to connect to past work*. To be explicit, these two measurements are *not* expected to agree.

Table 3.1 shows the results of applying Equation 3.1 to heat pulses generated by partial and full sawteeth from five representative Alcator C-Mod discharges, including two time periods from one of the discharges. Partial and full sawteeth were identified in these discharges manually, based on the criteria illustrated in Figures 3-1 and 3-2. As expected, the two measurements do not agree, with the ballistic transport

Shot	Confinement	$\chi_{Partial}^{pert} (m^2/s)$	$\chi_{Full}^{pert} (m^2/s)$
1120221011	L-Mode	1.1 ± 0.3	3.3 ± 0.5
1120626023	Ohmic (LOC)	2.7 ± 0.5	3.8 ± 0.7
1120626028	Ohmic (SOC)	1.7 ± 0.4	2.8 ± 0.5
1101209029	L-Mode	1.7 ± 0.3	4.7 ± 1.1
1101209029	I-Mode	2.0 ± 0.4	3.0 ± 0.9
1120221012	L-Mode	1.6 ± 0.4	3.6 ± 0.6

Table 3.1: Comparison of experimental perturbative thermal diffusivity as measured using Equation 3.1 on heat pulses generated by partial and full sawteeth. Data is taken from a variety of Alcator C-Mod discharges, including two time periods on 1101209029. All values are radially averaged. As expected these values differ, with the non-diffusive ballistic transport associated with full sawteeth causing an artificial increase in the measured diffusivity. Table modified from Reference [55].

associated with full sawteeth causing those values to exceed the values calculated with partial sawteeth by as much as a factor of 3. This quantitative observation is consistent with the qualitative observations regarding the shape of the heat pulses shown in Figure 3-2.

Again, measurements made with partial and full sawteeth are not expected to agree. The data shown here is further evidence that full sawteeth lead to non-diffusive ballistic transport, and that one should instead use partial sawteeth to calculate the perturbative thermal diffusivity.

On the other hand, a widely accepted means of measuring of the perturbative thermal diffusivity of a plasma is to analyze heat pulses generated by modulating applied electron cyclotron heating (ECH) [53, 56, 59, 86, 87]. If both modulated ECH and partial sawtooth-generated heat pulses are valid means of measuring the perturbative thermal diffusivity in a tokamak plasma, then these two measurements should agree to within experimental uncertainty. This work will show that these two measurements do indeed agree within uncertainty, at least on a limited set of ASDEX Upgrade discharges. This agreement provides strong evidence that partial sawteeth

are a valid method of measuring perturbative thermal diffusivity. This comparison was originally presented in Reference [64].

The modulated ECH measurements that are used in the comparison shown here are originally described in Reference [88]. While the original purpose of the modulated ECH measurements were not for comparison with partial sawteeth, data was published for two ASDEX Upgrade discharges, 31369 and 30693. The first of these 31369 conveniently also contains partial sawteeth, allowing application of the method described here. Unfortunately, 30693 did not have a sufficient number of partial sawteeth to make a reasonable measurement of the perturbative thermal diffusivity. Fortunately, however, the very next discharge, 30694, was an exact repeat of 30693 and did have a sufficient number of partial sawteeth to measure the perturbative thermal diffusivity. For this reason, one comparison of modulated ECH and partial sawteeth will be for discharge 31369, while the other will be for two repeat discharges, 30693 and 30694.

While this work focuses primarily on the partial sawtooth method of measuring perturbative thermal diffusivity, the fact that it is being compared to another method merits a brief discussion of how the modulated ECH method works. Instead of using heat pulses generated by naturally occurring partial sawteeth, modulated ECH generates heat pulses in the tokamak plasma by repeatedly turning electron cyclotron heating on and off at a relatively high frequency (100 Hz in the discharges in Reference [88], for example) [53, 56, 59, 86, 87]. Typically the modulated heating is deposited off axis to avoid interactions with sawtooth oscillations, for example at $\rho_{tor} = 0.46$ in the discharges described here. This modulation must be done for long enough to obtain good statistics on the pulse propagation, especially since the pulses tend to be very small in amplitude and diagnostic noise can be a problem. Typically the perturbative thermal diffusivity is calculated with the Fourier method [85] (mentioned earlier) instead of the Extended-Time-to-Peak method used here, since

	Modulated ECH		Partial Sawteeth	
Shot	Radial Range (ρ_{tor})	χ_{ECH}^{pert} (m^2/s)	Radial Range (ρ_{tor})	χ_{PST}^{pert} (m^2/s)
30693/4	0.34 - 0.44	2.2 ± 0.3	0.22 - 0.29	2.09 ± 0.62
31369	0.34 - 0.44	2.4 ± 0.9	0.23 - 0.29	2.87 ± 0.71

Table 3.2: Perturbative electron thermal diffusivity measured with modulated ECH heat pulses and partial sawtooth crash-generated heat pulses. Table modified from Reference [64], and modulated ECH data from Reference [88].

there is a large number of regularly spaced pulses, but as stated above, these two calculation methods have been shown to agree within experimental uncertainty [60]. More details of the analysis technique can be found in Reference [88], among others.

The results of the comparison of partial sawteeth and modulated ECH in ASDEX Upgrade discharges 31369 and 30963/4 are shown in Table 3.2. One caveat with this comparison is that while the two measurements were made in very similar radial locations, they do not quite overlap. The partial sawtooth heat pulses were measured in the radial range of $\rho_{tor} \approx 0.22 - 0.29$, while the modulated ECH measurements were made in the range of $\rho_{tor} \approx 0.34 - 0.44$. Again, while these do not overlap, the gap between the two measurements is approximately $0.05\rho_{tor}$, or 3 cm in physical space. The partial sawtooth measurement was constrained by the mixing radius at low radius and the pulse amplitude at high radius. The modulated ECH measurement was constrained by the ECH deposition location and (perhaps ironically) by the region in which sawtooth interference was too large [88].

Despite this caveat, Table 3.2 shows that measurements with modulated ECH and partial sawteeth agree to within experimental uncertainty in both cases. While the perturbative thermal diffusivity likely does vary with radius (since turbulent transport in general varies with radius), the gap between the two measurement locations is small enough that the variation is expected to be small. The two measurement method

also both capture the same qualitative trend, showing a higher perturbative thermal diffusivity in 31369 than in 30693/4.

The good agreement between the partial sawtooth method of measuring the perturbative thermal diffusivity laid out in this work and the more widely accepted modulated ECH method is strong evidence that both are measuring the same plasma property. Partial sawteeth will be used throughout this work to measure the perturbative thermal diffusivity of many plasmas both to search for experimental trends and to validate turbulent transport models of various kinds. The discharges shown below highlight one of the major advantages of the partial sawtooth method of measuring the perturbative diffusivity, as the measurement was obtained essentially ‘for free,’ without any active perturbation of the plasma. In addition, Alcator C-Mod does not have electron cyclotron heating (due to the difficulty of constructing high power sources at the frequencies required for Alcator C-Mod’s high magnetic field), and so partial sawteeth are the only viable method with which to measure the perturbative diffusivity.

This completes the discussion of how partial sawtooth crash-generated heat pulses are used to measure the perturbative thermal diffusivity of a tokamak plasma. The next section will describe a variety of experimental observations based on the perturbative thermal diffusivity on both Alcator C-Mod and ASDEX Upgrade. These observations are presented independently from any particular turbulent transport model validation efforts, which are the focus of Chapters 5 and 6.

3.4 Experimental Observations

Having developed the method of measuring the plasma perturbative thermal diffusivity via partial sawtooth-generated heat pulses, and confirming that this method agrees with the more established modulated ECH method, one may apply the method to a variety of plasma discharges. This section presents data from more than 50 Al-

cator C-Mod discharges and more than 20 ASDEX Upgrade discharges. The Alcator C-Mod data shows that I-mode plasmas tend to have a larger perturbative diffusivity than L-mode plasmas. In addition, trends of the perturbative diffusivity are observed with various plasma parameters on Alcator C-Mod, such as temperature, density, and their associated gradients. Data from both Alcator C-Mod and ASDEX Upgrade shows a strong correlation with two definitions of collisionality, indicating that it is possible to identify cross-machine empirical trends for the perturbative diffusivity.

3.4.1 Alcator C-Mod

The data from Alcator C-Mod presented here was originally published in Reference [55], and is taken from 56 discharges in which the plasma transition from a steady time period of L-mode to a steady time period of I-mode operation. The perturbative diffusivity stated for each mode of operation is calculated using partial sawtooth heat pulses that occurred during that mode of operation. Pulses that occurred during the transition from one mode to another were not analyzed. Some discharges had multiple distinct L-mode time periods (a steady L-mode, transition to I-mode, then a transition back to a steady L-mode), so a total of 79 points will be shown in subsequent figures describing the differences between L- and I-mode.

I-mode is a high performance regime [89, 90, 91] that has primarily been studied on Alcator C-Mod, but has also been observed on ASDEX Upgrade [92] and DIII-D [93]. I-mode plasmas are characterized by high energy confinement, similar to that found in a typical H-mode regime, but by a particle confinement that is more similar to that found in L-mode [89]. This is the result of I-mode plasmas having a temperature pedestal, but no density pedestal. I-modes also have a natural absence of Edge Localized Modes (ELMs), which are common to many H-mode plasmas [89]. The lower particle confinement in I-mode also avoids the core impurity accumulation that can cause disruptions in ELM-free H-mode plasmas. I-mode operation thus

exhibits many attractive attributes as a possible reactor regime, motivating further study of both its attributes and its operational space.

All of the discharges analyzed in this section were run with the ion $B \times \nabla B$ drift away from the active X-point (unfavorable ∇B drift for H-mode operation), since this generally allows more robust access to the I-mode regime [90]. In addition, all discharges were run with an on-axis toroidal magnetic field of 5.4 T, as higher magnetic field has been observed to lead to a higher threshold for the I- to H-mode transition, leading to a wider I-mode power window [94]. These discharges had line averaged densities ranging between approximately 0.6 to $2.1 \times 10^{20} \text{m}^{-3}$ and total plasma currents between 0.9 and 1.3 MA. All of the discharges were deuterium plasmas heated with minority hydrogen ion cyclotron resonance heating (ICRH).

This set of discharges was used to investigate two properties of the perturbative thermal diffusivity on Alcator C-Mod. First, the data set was utilized in order to generically observe the difference in perturbative diffusivity between L-mode and I-mode time periods within the same discharge. This is motivated by past gyrokinetic simulations with the gyrokinetic code GYRO [95], which predicted that I-mode plasmas on Alcator C-Mod would have a higher perturbative diffusivity than similar L-mode plasmas [96, 97]. In addition, these measurements were utilized to determine the correlation between perturbative diffusivity and many other plasma parameters, such as density, temperature, their gradients, etc. Each of these two investigations will now be described in further detail.

Experimental Perturbative Thermal Diffusivity in L- and I-mode

First, consider the comparison between L-mode and I-mode within the same discharge. The choice to compare discharges only with themselves, just at different times, was made in order to minimize the variation of as many plasma parameters as possible. In general, the L- and I-mode time periods within each discharge had

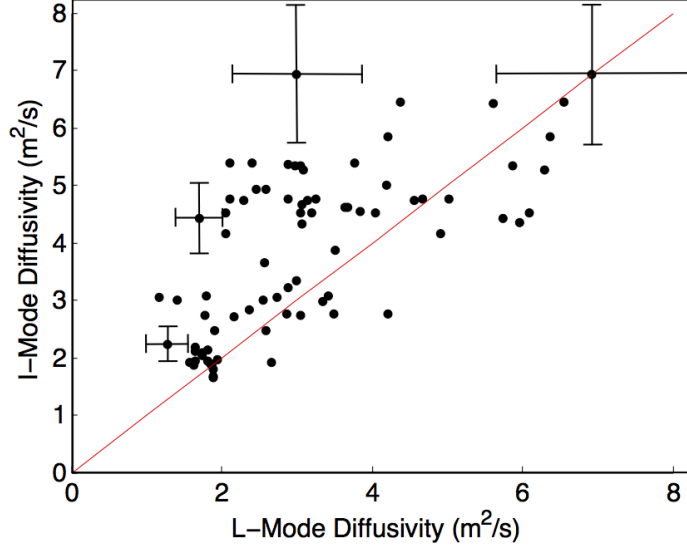


Figure 3-4: Comparison of the perturbative thermal diffusivity calculated via partial sawteeth in L- and I-mode time periods from within the same discharge. Each point is data from one discharge, and the x- and y-axes are the L-mode and I-mode perturbative diffusivities, respectively. Discharges with multiple distinct L-mode time periods are represented by multiple points. The solid red line indicates equal diffusivities in both time periods. Representative uncertainties are given. Figure from Reference [55].

the same current, the same plasma shape (triangularity, elongation, etc.), and the same magnetic field. The transition from L- to I-mode was generally initiated by an increase in ICRH power. The density was kept as steady as possible over this transition, though often the increase in RF (radio frequency) power led to a slight increase in the plasma density, as is often seen on Alcator C-Mod. The increase in heating power and the transition from L- to I-mode always led to an increase in both ion and electron temperatures, which is to be expected in the transition from a low to a high confinement regime.

With these caveats in mind, the results of the comparison between L- and I-mode are shown in Figure 3-4. Each point on this plot represents a single comparison between an L-mode time period and an I-mode time period within one discharge. The x-axis is the perturbative diffusivity calculated from partial sawteeth in the L-mode portion of the discharge, and the y-axis is the perturbative diffusivity calculated from

partial sawteeth in the I-mode portion of the discharge. If there were multiple distinct L-mode time periods within a single discharge, two points are shown (comparing to the same I-mode time period). The solid red line indicates equal perturbative diffusivity in L- and I-mode. Points above the line represent a discharge in which the perturbative diffusivity was higher in the I-mode time period than in the L-mode time period, and vice versa. Representative uncertainties (calculated with Equation 3.8) are shown.

While there is a considerable amount of scatter in the data shown in Figure 3-4, and many of the points lie within uncertainty of the line of equal perturbative diffusivities, overall there are more points above the red line than below it, indicating that in general the perturbative diffusivity is higher in the I-mode time period of a given discharge than in the L-mode time period. Note, however, that such a difference may also be due to differences in the plasma temperature and temperature gradient during the L- and I-mode time periods, as these parameters are shown to also be correlated with the perturbative diffusivity later in this section.

While the comparison shown in Figure 3-4 may not reveal any definitive conclusions, it is consistent with past gyrokinetic predictions that I-mode plasmas should have a higher perturbative thermal diffusivity than similar L-mode plasmas [96, 97]. In addition, these experimental observations are relevant to the gyrokinetic validation work that will later be presented in Chapter 5, since part of that work focuses on validating the gyrokinetic code GYRO in L- and I-mode plasmas on Alcator C-Mod. The particular discharge chosen for those simulations will be described in further detail in Chapter 5.

Correlation Between Experimental Perturbative Diffusivity and Plasma Parameters

Consider now the relationship between the perturbative thermal diffusivity and other plasma parameters, such as the electron temperature and density, and their gradients, both at the location of the perturbative diffusivity measurement and in the core of the plasma. This section presents correlations between the measured perturbative diffusivity and a number of other global and local plasma parameters for the discharges considered on Alcator C-Mod. It does not, however, attempt to comment on the physical mechanisms behind these relationships. This section also compares the perturbative diffusivity to the amplitude of the partial sawtooth-generated heat pulse, in order to ensure that this amplitude is uncorrelated with the measurement, as such a correlation would indicate that the measurement method is somehow influenced by the size of the partial sawtooth crash. Section 3.4.2 will expand these results to include data from ASDEX Upgrade.

While comparisons with many different plasma parameters were considered, only those that showed meaningful correlations are explained in detail here. Starting with global parameters, central electron temperature, RF heating power, and stored energy seemed to have little to no correlation with the perturbative diffusivity.

Core line averaged density, \bar{n}_e , and plasma current, I_p , on the other hand, both showed a fairly clear correlation with the perturbative diffusivity. Experimentally, however, current and density are themselves correlated in tokamak operation (higher current allows stable access to higher densities by increasing the Greenwald density limit), so it was important to determine if both of these parameters correlated separately with perturbative diffusivity. To this end, correlations between perturbative diffusivity and density at fixed current, and current at fixed density were calculated. This analysis revealed that current at fixed density does not show a trend with perturbative diffusivity, while density at fixed current does, indicating that current only

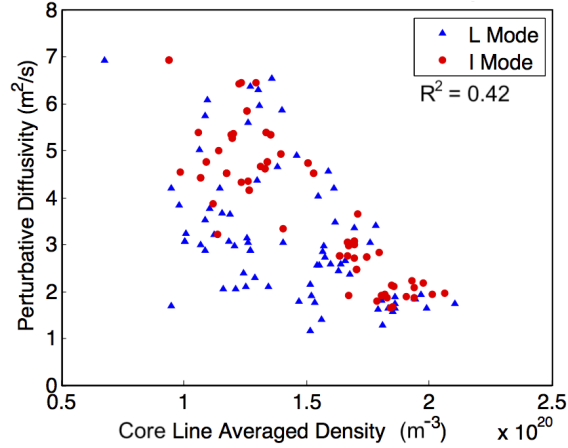


Figure 3-5: Perturbative thermal diffusivity plotted against the line averaged density of both L- and I-mode plasmas. Each point represents either an L-mode or an I-mode time period of a given shot. Blue triangles represent L-mode and red circles represent I-mode phases. Figure originally from [55].

correlated with perturbative diffusivity through its connection with density.

A plot of perturbative diffusivity against core line-averaged density is shown in Figure 3-5, using data from all currents. Each point in this plot represents one time period of either L- or I-mode. Blue triangles represents L-mode plasmas and red circles represent I-mode plasmas. This figure reveals a clear trend of decreasing perturbative diffusivity as the plasma moves toward higher density. There also seems to be a smaller range of perturbative diffusivity as the density increases. In other words, it appears as if the maximum perturbative diffusivity decreases as one increases the density, but that the minimum perturbative diffusivity stays roughly constant.

More quantitatively, one can use a linear regression to calculate the p-Value and R^2 value for the data. The p-Value is the probability of no correlation given the data (the null hypothesis), and the R^2 value measures the amount of variation in the perturbative diffusivity that can be explained by the variation in the line-averaged density [98]. A linear regression of the perturbative diffusivity with the line-averaged density data shown in Figure 3-5 gives a p-Value of 5.0×10^{-20} and an R^2 of 0.42, indicating a reasonable correlation. These values will be compared to the values

calculated for other plasma parameters in order to determine which correlation is strongest.

Consider now the local plasma temperature and density, and their gradients. Figure 3-6 shows the perturbative diffusivity from the same set of discharges plotted against local density, density gradient ∇n_e , and normalized density gradient scale length a/L_n , (where a is the minor radius and $L_n = n_e/\nabla n_e$). These values are averaged over the radial region in which the perturbative diffusivity was measured. Perhaps unsurprisingly, the local density shows an almost identical trend with perturbative diffusivity as did the line-averaged density. In this case, a linear regression gives a p-Value of 5.3×10^{-14} and an R^2 of 0.31, which are similar, though not quite as conclusive as the line-averaged density. Again, the local density seems to correlate to a maximum perturbative diffusivity, though not to a minimum.

The density gradient, in Figure 3-6 (b), shows a similar negative correlation with the perturbative diffusivity, though the shape of this correlation seems to have more structure. Instead of a decreasing maximum value, an increasing density gradient correlates to a cutoff value, at around $5 \times 10^{20} \text{m}^{-3}/\text{m}$, below which there is a linearly decreasing perturbative diffusivity and above which there is a roughly constant range of perturbative diffusivities. There are, however, a few points that defy this trend. Even though it would appear that a linear fit may not be entirely appropriate, for continuity with the other trends presented here one was performed, giving a p-Value of 6.4×10^{-9} and an R^2 of 0.20. In addition, since the normalized density gradient scale length, a/L_n , is often associated with turbulence drive terms [99], Figure 3-6 (c) shows the perturbative thermal diffusivity plotted against a/L_n . Linear regression, however, revealed little to no correlation.

Consider now the local plasma temperature and its gradient. Figure 3-6 (d) shows the perturbative diffusivity plotted against local temperature, revealing a positive correlation. A linear regression of this data gives a p-Value of 7.0×10^{-10} and an

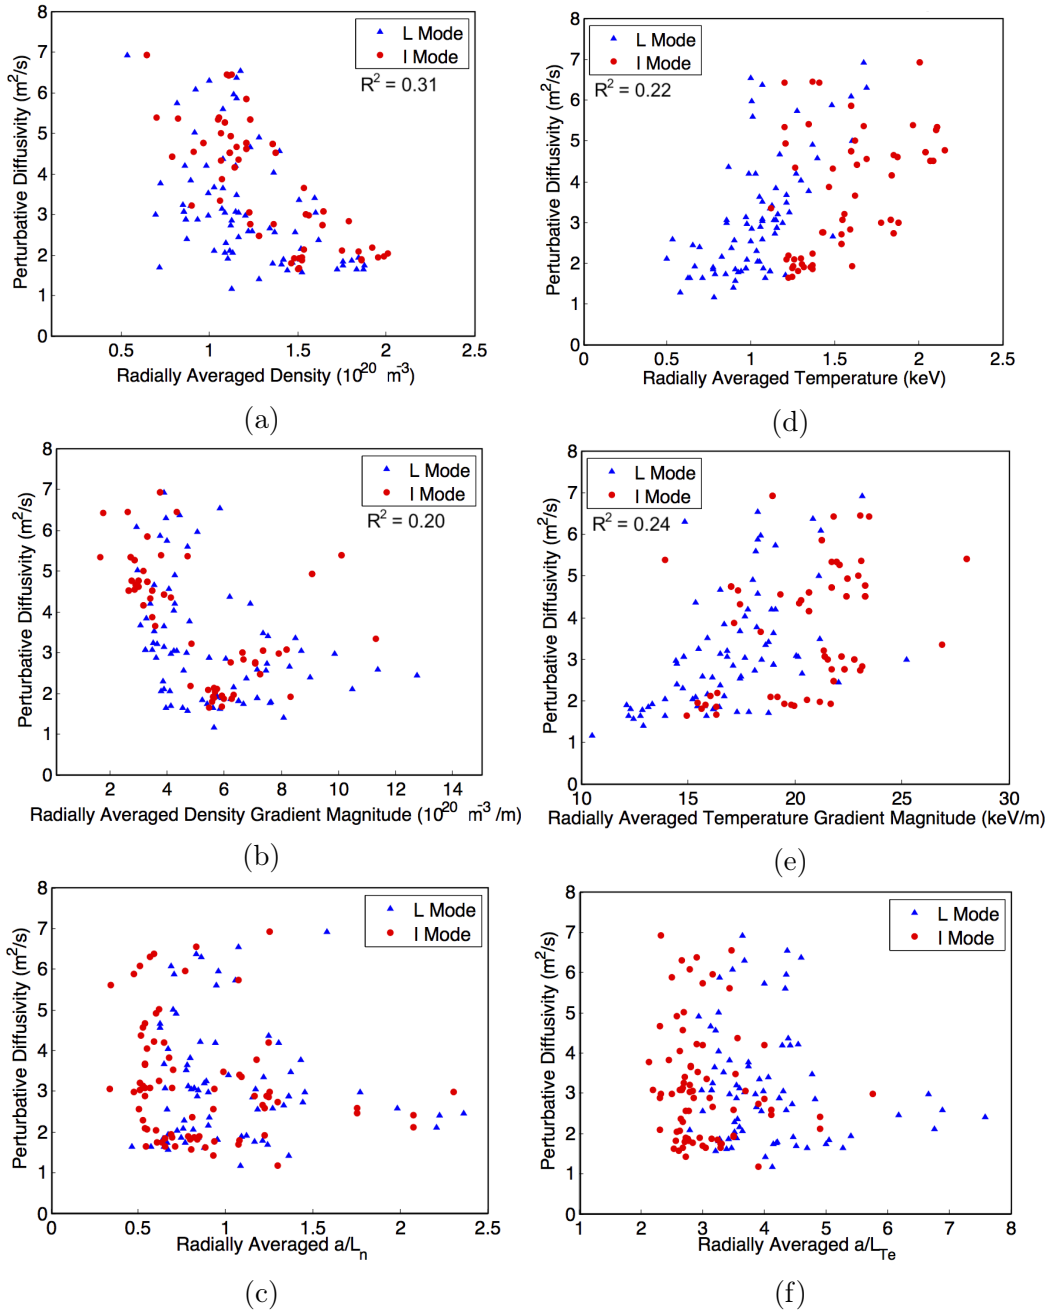


Figure 3-6: Perturbative thermal diffusivity plotted against (a) electron density, (b) density gradient, (c) a/L_n , (d) electron temperature, (e) temperature gradient, and (f) a/L_{Te} , all averaged over the radial range of the diffusivity measurement. Symbols are used as above. Figure originally from [55].




Figure 3-7: Perturbative thermal diffusivity normalized by the Gyro-Bohm factor of $T^{3/2}$ plotted against (a) a/L_{Te} and (b) a/L_n . Symbols used as above. Figure originally from [55].

R^2 of 0.22. A similar trend is observed with the temperature gradient, in Figure 3-6 (e), giving a p-Value of 1.3×10^{-10} and an R^2 of 0.24. Note that in both of these cases, the L- and I-mode plasmas are somewhat separated on the plots, as I-mode in general tends to have higher temperatures and temperature gradients than L-mode. Interestingly, when one plots the perturbative thermal diffusivity against the normalized electron temperature gradient scale length, a/L_{Te} , as in Figure 3-6 (f), one observes little to no linear correlation. One does, however, observe that in general, I-modes tend to have lower values of a/L_{Te} than L-modes. This is in fact consistent with other observations that I-mode tends to have lower temperature fluctuation levels than L-mode [100], since once again one associates a/L_{Te} with turbulence drive.

Having analyzed both temperature and density separately, one may ask about correlations with combinations of these local parameters. This was indeed investigated, and the results of this analysis, in particular correlations with various definitions of collisionality, are discussed in the next subsection.

First, however, consider plotting the perturbative diffusivity normalized by the local temperature to the 3/2 power against a/L_{Te} and a/L_n , as is done in Figure 3-7. This normalization was chosen as $T^{3/2}$ is commonly associated with the Gyro-Bohm

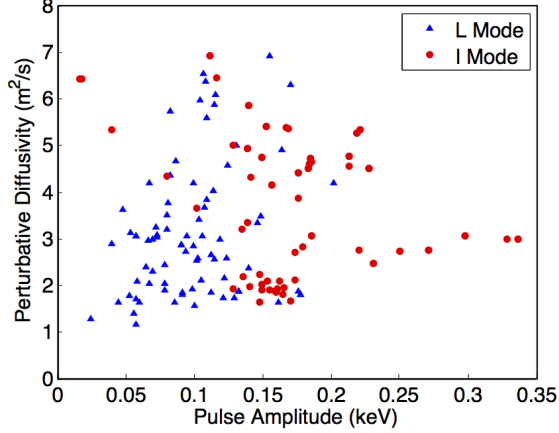


Figure 3-8: Perturbative thermal diffusivity plotted against the partial sawtooth-generated heat pulse amplitude, averaged over the radial range of the diffusivity measurement. Symbols used as above. Figure originally from [55].

factor (a theoretical scaling for turbulent transport), and normalization by this factor is common in other perturbative transport studies [101]. While figure 3-7 does not reveal any particularly interesting trends, it does show that this normalization tends to group the I-mode plasmas closer together, while the L-modes remain fairly spread out. One possible explanation for this is that the I-modes tend to be hotter than the L-modes, which would more tightly group the discharges when normalized inversely with temperature. Returning to linear regressions, the plot against a/L_{Te} gives a p-Value of 2.0×10^{-21} and an R^2 of 0.45, and the plot against a/L_n gives a p-Value of 3.0×10^{-10} and an R^2 of 0.13.

Before moving on to the collisionality data from two machines, a final check was made to confirm that the amplitude of the partial-sawtooth generated heat pulse did not correlate with the measured perturbative diffusivity. If it were the case that the pulse amplitude and measured diffusivity were correlated, one may postulate that one is no longer observing a purely diffusive process, and that the measurement is then somehow contaminated with other non-diffusive effects. Figure 3-8 reveals, however, that there is little to no correlation between the measured perturbative thermal diffusivity and the pulse amplitude. This is further evidence that the partial sawtooth

heat pulses are not contaminated with the ballistic transport that complicates the use of full sawtooth heat pulses in similar analysis, and that these measurements are therefore properly measuring the plasma perturbative diffusivity.

3.4.2 Cross-Machine Correlations on ASDEX Upgrade and Alcator C-Mod

In addition to the correlations between perturbative thermal diffusivity and various plasma parameters observed on Alcator C-Mod, combined data from both Alcator C-Mod and ASDEX Upgrade shows trends with various definitions of collisionality, as originally described in Reference [64]. In particular, data from the L-mode time periods of the Alcator C-Mod discharges just shown will be combined with data from 24 L-mode discharges on ASDEX Upgrade.

First, all of the dependencies just described (temperature, density, gradients, etc.) were also investigated on ASDEX Upgrade, finding similar trends to those seen on Alcator C-Mod. In addition, the availability of data from two fairly different machines allowed for the exploration of new parametric dependencies over a wider range, such as the plasma collisionality. While there are many definitions of collisionality, which generally normalize one of the plasma collision frequencies to some other frequency in the plasma (bounce frequency, cyclotron frequency, etc.), two in particular were chosen for investigation here.

Consider first a collisionality, ν^* , based on normalization of the electron-ion collision frequency by the electron bounce frequency [6]:

$$\nu^* \approx 0.01 \left(\frac{R_0}{r} \right)^{3/2} \left(\frac{qR_0 n_{20}}{T_k^2} \right) \quad (3.9)$$

where R_0 is the major radius, r is the local minor radius, q is the safety factor, n_{20} is the plasma density in units of 10^{20}m^{-3} and T_k is the electron temperature in

units of keV.

As seen in Figure 3-9 (a), this definition of collisionality tends to bring the data from both machines into the same range, since Alcator C-Mod is smaller than ASDEX Upgrade, but tends to have much higher density. The perturbative diffusivities from both machines overlay quite nicely, with increasing ν^* generally leading to decreasing perturbative diffusivity. The trend observed in collisionality is consistent with the trends observed in density and temperature in the previous section, since increasing density and decreasing temperature led to decreased perturbative diffusivity, and ν^* goes as density over temperature. On the log-log axes in Figure 3-9 (a) the perturbative diffusivity and ν^* are correlated almost linearly, indicating an exponential relationship. More quantitatively, a fit of $(\nu^*)^{-0.5}$ has an unweighted R^2 value of 0.53, which exceeds all of the R^2 values from the previous section, despite including data from two different machines.




Figure 3-9: Plots of perturbative diffusivity from 24 ASDEX Upgrade discharges (red) and 56 Alcator C-Mod discharges (blue) against two definitions of collisionality. All discharges are L-mode plasmas. (a) shows ν^* , the electron-ion collision frequency normalized by the bounce frequency, and (b) shows a normalized collisionality used in the gyrokinetic code GENE [26]. The point circled in green in (b) is the discharge used in the GENE validation study in Chapter 5. Figures originally presented in Reference [64].

Consider next an alternative definition of the plasma collisionality, such as the following species-independent normalized collision frequency used in the gyrokinetic

code GENE [26]:

$$\nu_{GENE} = 2.3031 \cdot 10^{-5} \frac{a \cdot n_{19} \cdot \ln\Lambda}{T_k^2} \quad (3.10)$$

where a is the square root of the edge toroidal flux divided by π times the reference toroidal magnetic field, $a = \sqrt{\Phi_{edge}/\pi B_{ref}}$ (generally equal to the machine minor radius except in cases of extreme Shafranov shift) [102], n_{19} is the plasma density in units of 10^{19}m^{-3} , T_k is the electron temperature in units of keV, and $\ln\Lambda$ is the Coulomb Logarithm. Internally in GENE this quantity is labelled ν_c , but here it will be called ν_{GENE} for clarity. Note that GENE does use species-dependent collision frequencies when performing relevant calculations, but that this species-independent normalized collisionality can be used to parametrize the plasma in question [26].

The perturbative diffusivity is plotted against ν_{GENE} in Figure 3-9 (b). With this species-independent definition of collisionality the data from Alcator C-Mod and ASDEX Upgrade are somewhat more separated than with ν^* , leaving some regions on either end of the plot with data from only one machine. As with ν^* , however, there is a very clear trend of decreasing perturbative diffusivity with increasing collisionality, though the exact fit is slightly different. In particular, a fit of $0.15(\nu_{GENE})^{-0.44}$ gives an R^2 value of 0.52, which is nearly identical to that for the fit to ν^* . Note that one of the ASDEX Upgrade discharges shown here, that circled in green in Figure 3-9 (b), is the same discharge that is used for a validation study of GENE in Chapter 5.

The physical means by which the collisionality may influence the perturbative diffusivity are likely complex, but may be related to the relative strength of the ITG (ion temperature gradient) and TEM (trapped electron mode) turbulent modes in these plasmas. Reference [103] notes that lower collisionality tends to favor stronger destabilization of TEM compared to ITG, while higher collisionality tends to increase the relative strength of ITG modes. Furthermore, Reference [104] shows that increased TEM drive correlates with increased perturbative diffusivity (as measured with mod-

ulated ECH). This may be unsurprising, since TEM is in part driven by the electron temperature gradient while the ITG is driven by the ion temperature gradient, and the perturbative diffusivity measures the relationship between electron heat flux and the *electron* temperature gradient. The experimental correlations presented here from Alcator C-Mod and ASDEX Upgrade are consistent with this interpretation.

The results presented here establish that the partial sawtooth method of measuring the perturbative diffusivity can be applied to many discharges on both Alcator C-Mod and ASDEX Upgrade, and reveal that the physics of the perturbative diffusivity are not machine dependent. A more detailed investigation of the physical underpinnings of the relationship between the perturbative diffusivity and various plasma parameters may be the subject of future work.

Conclusions

In conclusion, this chapter has described a method of measuring the perturbative thermal diffusivity in a tokamak plasma by tracking the propagation of heat pulses generated by partial sawtooth crashes. First, the characteristics of a partial sawtooth crash were described, including what distinguishes partial and full sawteeth. Next, details of how the perturbative diffusivity is calculated based on the partial sawtooth-generated heat pulses were given, including a comparison of two equivalent calculation procedures. The results of these calculations were then compared to similar calculations based on full sawteeth (to show that the partial sawtooth method is superior) and to measurements based on the more established modulated electron cyclotron heating method (to show that the partial sawtooth and modulated ECH methods agree). Finally, a variety of experimental observations from Alcator C-Mod and ASDEX Upgrade were presented, focusing on comparisons between L- and I-mode plasmas and the correlation between perturbative diffusivity and various plasma parameters. The strongest of these correlations was between the

perturbative diffusivity and various definitions of collisionality, using data from both Alcator C-Mod and ASDEX Upgrade.

The method of measuring perturbative diffusivity described in this chapter was primarily developed as a tool for the validation of turbulent transport models. In addition to measurement of electron temperature fluctuations obtained using Correlation Electron Cyclotron Emission (CECE) diagnostics, as described in Chapter 4, measurements of perturbative diffusivity are one of the key validation metrics used in Chapter 5 to validate the gyrokinetic codes GYRO and GENE and Chapter 6 to validate the gyro-fluid code TGLF. Further details of how this validation was performed is given in each respective chapter.

Chapter 4

Electron Temperature Fluctuations from Correlation Electron Cyclotron Emission

This chapter describes the Correlation Electron Cyclotron Emission (CECE) diagnostic, which is used to measure electron temperature fluctuations, \tilde{T}_e/T_e , in tokamaks, the second validation constraint considered in this thesis. Chapter 2 described the theoretical basis of the CECE diagnostic, focusing on the emission process. This chapter details the practical implementation of such a diagnostic, including the hardware that was built on both Alcator C-Mod and ASDEX Upgrade as part of this thesis. In addition, new data analysis techniques were developed in order to more accurately extract temperature fluctuation spectra and total fluctuation levels from the raw CECE data. Finally, this chapter presents experimental CECE observations from both Alcator C-Mod and ASDEX Upgrade, focusing on experimental trends. As with the perturbative diffusivity from Chapter 3, CECE data will also be used as a validation constraint in Chapters 5 and 6.

4.1 CECE Radiometers

All of the CECE diagnostics described in this work are based on heterodyne radiometers. Generically, a heterodyne radiometer works as follows. First, the electron cyclotron emission is collected from the plasma using some form of antenna. This signal, which includes a broad range of frequencies, is then band-pass filtered with a fairly wide passband (roughly 10 GHz) in order to narrow the signal to only the frequencies that are relevant to the ECE spectrum emitted within the plasma volume. This signal, which is typically in the hundreds of GHz, is then down-converted to an intermediate frequency, in the ones or tens of GHz, by mixing with a local oscillator (LO). The intermediate frequency signal is divided into several channels, which are each bandpass filtered at different center frequencies to generate signals that are well localized in space (due to the spatial variation of the magnetic field and thus the ECE frequency in the tokamak volume). Each localized signal is then digitized, and analyzed digitally.

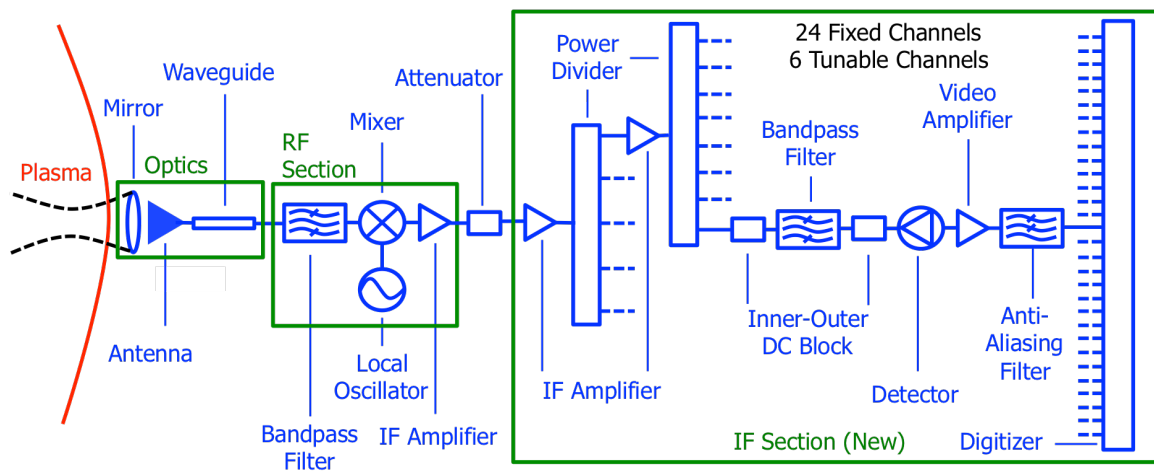


Figure 4-1: Schematic block diagram of the new ASDEX Upgrade CECE system. There are twenty four fixed frequency bandpass channels and six tunable bandpass channels, for a total of thirty channels. Figure from Reference [81].

Figure 4-1 shows a block diagram for the CECE radiometer on ASDEX Upgrade. The figure shows all of the components described above, as well as some additional

electronic components that are required for the operation of the system.

As described in Chapter 2, CECE measures temperature fluctuations by correlating the signals of two closely spaced, but not overlapping, radiometer channels. The radial dimension of the emission volume from which the electron cyclotron emission (ECE) is collected is set by the bandwidth of the radiometer filters, and is affected by the Doppler and relativistic shifts in the cyclotron frequency of the emitting electrons [7]. In the poloidal direction, the size of the emission volume is set by the local diameter of the optical beam that collects the ECE for the diagnostic. This beam spot size can be very important, as it generally sets the minimum resolvable fluctuation size, or alternatively, the maximum resolvable fluctuation wavenumber, $k = 2\pi/\lambda$, where λ is the wavelength of the fluctuation. For the purposes of CECE, the beam spot size, w_0 , used is the $1/e^2$ power (or equivalently the $1/e$ electric field) diameter. For comparison to gyrokinetic and other turbulent models, the important parameter will be the poloidal wavenumber times the ion gyro-radius (often evaluated with the electron temperature), typically expressed as $k_\theta \rho_s$, where $\rho_s = \sqrt{2T_e/m_i}$. This therefore depends on both the physical optics and the specific plasma under consideration.

This section will go on to describe the hardware on both Alcator C-Mod and ASDEX Upgrade, including the optical spot sizes.

4.1.1 Alcator C-Mod Hardware

A significant portion of the work on the Alcator C-Mod CECE hardware was performed as part of a prior graduate student's Ph.D. thesis [7], and is described further in References [75] and [76]. A description of the system is repeated here as it was used to collect all of the Alcator C-Mod data that is used later in this work. In addition, the intermediate frequency (IF) section of this diagnostic was rebuilt as part of this thesis work, with the intention of prototyping the ASDEX Upgrade system, described afterward, as well as reducing system noise.




Figure 4-2: CECE hardware on Alcator C-Mod. (a) In-vessel optics. (b) Re-built intermediate frequency (IF) section.

The ECE radiation is collected using a set of in-vessel optics pictured in Figure 4-2 (a). This optical system consists of a parabolic stainless steel focussing mirror, a flat mirror, a high density polyethylene (HDPE) collimating lens, and a corrugated, high gain scalar horn antenna. The system was designed to collect radiation from 230 to 270 GHz [105]. The exact spacing and focal lengths of the lenses and mirrors changed during the life of the system, as is detailed in Reference [7] (giving spot size diameters of between roughly 0.5 and 1.0 cm over the radial range of interest). This knowledge proved important for the work presented here, as the data from different discharges that were analyzed came from different eras of hardware setup. This system collected second harmonic, X-mode electron cyclotron emission from the plasma, and was located approximately 7 cm above the mid-plane of the low-field side of the machine.

The horn antenna output is fed into the radio frequency (RF) section of the CECE diagnostic. The signal was band-pass filtered with a pass band of 232 to 248 GHz and then mixed with a 250 GHz local oscillator, producing an IF signal of 2 - 18 GHz. At a standard Alcator C-Mod magnetic field of 5.4 T, this frequency range covers a plasma region of approximately ρ_{tor} from 0.65 to 0.90, where ρ_{tor} is the square root of the normalized toroidal magnetic flux. This lower frequency signal is then amplified with a 33 dB low noise amplifier before transmission along a 6.1 m low loss SMA cable to the IF section. The block diagram for the system on Alcator C-Mod is very similar to that for ASDEX Upgrade, shown in Figure 4-1, though with fewer channels.

The original IF section and a previous upgrade are described in References [75] and [76]. The rebuild done as part of this thesis is pictured in Figure 4-2 (b). This most recent incarnation of the IF section consisted of a 30 dB low noise 2-18 GHz amplifier, a four-way power splitter, DC blocks, bandpass filters (fixed-frequency and tunable YIG), DC blocks, Schottky diode detectors, high-pass filters to remove the DC component (background temperature), 0 - 6 MHz video amplifiers, and 1MHz low-pass anti-aliasing filters. As can be seen in the picture, care was taken to minimize cable lengths, to the point that the chassis contained only two SMA cables, with all other connections accomplished with SMA adaptors.

The two fixed frequency filters were 100 MHz bandwidth cavity filters, with center frequencies of 4.00 and 4.14 GHz, for a center-to-center separation of 140 MHz. The two Yttrium Iron Garnet (YIG) filters had a bandwidth of 100 MHz and center frequencies that could be tuned between 6 and 18 GHz. The tuning was accomplished via an externally applied current, with currents of 0.3 A to 0.9 A corresponding approximately linearly to the frequencies between 6 and 18 GHz. The external current was applied with a remotely controlled, two-output, tabletop power supply.

Finally, the signal was digitized at 10 MS/s with a 16 channel, 14 bit D-TACQ digitizer, and then stored in the Alcator C-Mod MDSplus data storage system.

4.1.2 ASDEX Upgrade Hardware

The CECE system on ASDEX Upgrade was intentionally constructed so as to be able to measure both electron temperature fluctuations and the phase between temperature and density fluctuations (n-T phase). The CECE system is not itself able to measure density fluctuations, but in concert with a reflectometer system can be used to measure the phase angle between the two channels of fluctuations. This work will primarily focus on the CECE aspect of the system, not on the n-T phase capabilities, but the dual purpose does impact the design of the hardware.

One such impact is on the system optics. The ASDEX Upgrade CECE system shares all of its front-end optical components with the existing steerable Doppler reflectometer on ASDEX Upgrade [106]. This optical system includes a steerable elliptical mirror, a smooth-bore Gaussian-beam antenna, a 38 mm diameter oversized waveguide, two waveguide tapers, and fundamental waveguide. This optical setup gives a beam diameter of approximately 3.0 to 3.5 cm in the outer core of the plasma [107]. The signal is split between the two systems while in oversized waveguide using a 3dB grid splitter. See Reference [80] for further information regarding the optical components and the interaction between the two systems.

The fundamental waveguide from the optical system connects to the radio frequency (RF) section, the internal components of which are picture in Figure 4-3 (a). Interchangeable sideband filters on this section allow the measurement of frequencies of either 105 to 113 GHz or 117 to 125 GHz. For a typical ASDEX Upgrade discharge with an on-axis magnetic field of 2.5 T, this corresponds to a radial range of roughly $\rho_{tor} = 0.4$ to 1.0, where ρ_{tor} is the square root of the normalized toroidal magnetic flux. This particular RF section can also be interchanged with others in order to access completely different frequency ranges, for use on discharges with different field or to focus on particular regions of the plasma. As discussed above, the RF section





Figure 4-3: Original CECE system on ASDEX Upgrade. (a) Radio frequency (RF) section electronics, before being mounted inside of a protective box. From left to right: Gunn oscillator, multiplier, mixer, 2-18 GHz amplifier (top), sideband filter. (b) Original IF section with four fixed and two tunable channels.

then mixes the high frequency signal with the signal from a local oscillator (a Gunn oscillator and multiplier) to get an output of 2 to 14 GHz. The signal is then transmitted from the RF section to the intermediate frequency (IF) section through an approximately 30 cm long SMA cable.

Both the original and upgraded IF sections on ASDEX Upgrade were constructed as part of this thesis work. The original IF section on ASDEX Upgrade is shown in Figure 4-3 (b). This section consisted of the following: a 2-18 GHz amplifier, a two way power divider, a two-way and four-way power divider (for a total of six channels), DC blocks, four fixed frequency and two tunable YIG bandpass filters, DC blocks, Schottky diode detectors, high-pass filters, and 0 - 6 MHz video amplifiers. The two YIG filters were tunable between 6 and 18 GHz. The fixed frequency bandpass filters were interchangeable, with center frequencies of between 4 and 10 GHz, and bandwidths of either 100 or 200 MHz.


The upgraded IF section on ASDEX Upgrade is shown in Figure 4-4 (and was shown schematically in Figure 4-1) and is described in Reference [81]. The upgrade increased the total number of channels in the CECE system from six to thirty, with




(a)



(b)



(c)



(d)

Figure 4-4: Upgraded CECE IF section on ASDEX Upgrade. (a) Amplifier and divider chassis. (b) Fixed frequency chassis containing eight fixed frequency bandpass filters. The system contains three of these chassis. (c) YIG chassis. The system contains two of these chassis. (d) Four of the six CECE IF chassis and the digitizer installed at ASDEX Upgrade.

twenty-four fixed and six tunable channels. This increase in number of channels did, however, require more space, and so the new IF section consisted of six separate chassis, rather than a single chassis (though four of the chassis were only one unit tall

in a 19-inch rack).

The first chassis, called the amplifier-divider chassis and pictured in Figure 4-4 (a), contains a low-noise 30 dB 2-18 GHz amplifier (two are included, but only one is normally used), a two way power divider, a two-way and four-way power divider in parallel (for a total of six channels), and then six identical low-noise 30 dB 2-18 GHz amplifiers. Of the six outputs from the amplifier-divider chassis, only five are currently used: three go to fixed-frequency filter sections, and two to YIG filter sections. Each of the amplifiers is attached to a copper heat sink.

The fixed-frequency filter chassis, one of which is pictured in Figure 4-4 (b), each contain eight channels based on fixed-frequency cavity bandpass filters. More specifically, each chassis contains an eight-way power divider, DC blocks, eight cavity filters of various bandwidths and center frequencies, DC blocks, Schottky diode detectors, 0 - 6 MHz video amplifiers, and 2 MHz low-pass anti-aliasing filters (custom fitted inside the video amplifiers). Most of the individual components are the same type as used in the original system, though there are no high-pass filters. It was discovered that the digitizers had enough dynamic range to enable resolution of both the signal mean and the fluctuations around this mean. In addition, measuring the mean signal allows accurate characterization of the exact filter bandwidth, by comparing the signal mean and standard deviation with a known noise source as input. This measurement allows an independent check of the true bandwidth of the filter.

Fixed filters of different center frequencies and bandwidths can be changed out fairly easily, so the system is quite flexible, but in the standard configuration the system operates with 100 MHz bandwidth filters spaced by 125 MHz and 200 MHz bandwidth filters spaced by 250 MHz. These two sets of filters were arranged in a comb configuration to cover roughly the range from 4 to 8 GHz. The wider bandwidth filters allow for a wider radial range of simultaneous measurements, while the narrower bandwidth filters covered a portion of that range with finer resolution, allowing

for accurate assessment of the radial correlation length of the temperature fluctuations. In addition, including both sets with some overlapping center frequencies allows consistency checks of measurements made on entirely independent channels.

In addition to the three fixed-frequency filter chassis, there are two YIG filter chassis, one of which is pictured in Figure 4-4 (c). The Yttrium Iron Garnet (YIG) filters are bandpass filters of a fixed bandwidth but a variable center frequency. The particular filters used here have a bandwidth of 200 MHz and center frequencies which are tunable from 6 GHz to 14 GHz. The tuning is accomplished via an externally applied control voltage, from 0 to 10V, tuning the center frequency roughly linearly. The ASDEX Upgrade system is set up such that the control voltage is automatically recorded into the shotfile system whenever a discharge is taken. Other than the different filters, the YIG chassis have all of the same components as the fixed-frequency chassis (detectors, amplifiers, etc.). One chassis has 4 channels, and the other has 2 channels with space for 2 more in a possible future upgrade. Each YIG filter is attached to an aluminum heat sink.

Finally, the outputs from the five filter chassis are digitized in a pair of synchronized 16 channel, 14 bit digitizers, typically at 4 MS/s, though this can also be varied.

Four of the six IF chassis and the digitizer as they are installed are pictured in Figure 4-4 (d). The amplifier-divider chassis and one of the fixed-frequency chassis are mounted on an adjacent shelf in the rack. Both of the YIG chassis are actively cooled with a fan beneath them. All together, the system has 30 separate channels (24 fixed and 6 tunable) in 6 modular chassis (one amplifier-divider chassis, three fixed-frequency filter chassis, and two YIG chassis). Even though the new system significantly increased the number of channels over the original, optimization of the physical layout within each chassis kept the overall size relatively modest. This internal optimization also had the effect of reducing cable lengths and superfluous adapters within the chassis, reducing electronic noise and losses. Another careful considera-

tion was the ground isolation within each chassis, which also aided in reducing system noise. The green G10 (fiberglass laminate) boards seen in Figure 4-4 are part of this isolation scheme. In particular, care was taken to maintain separate signal and chassis grounds throughout the system, preventing power supply and ambient noise from affecting the signal. This seemed particularly important after multiple experiences of static shocks when working with the 19 inch rack that was available in the ASDEX Upgrade torus hall.

This concludes the description of the physical CECE hardware at ASDEX Upgrade.

4.2 Data Analysis Techniques

The data collected by the CECE hardware just described is related to the electron temperature fluctuation level via the theory described in Chapter 2 and statistical analysis techniques that will be described here. These techniques were originally published in Reference [81]. This work builds off of statistical analysis techniques from References [108, 109, 110, 111, 112, 113, 114, 115, 116, 117] and electron cyclotron theory presented in References [69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 80, 105, 118, 119]. The analysis techniques used to analyze the CECE data evolved somewhat through the course of the work presented here, having been adapted to be more general and more mathematically rigorous. This section will first present the final forms of the equations that were developed and presented in Reference [81], then will discuss earlier forms that were used and how these results differ from the final forms. For full derivations of these equations, see Appendix C.

4.2.1 Calculation of Temperature Fluctuation Levels

Most generally, the total temperature fluctuation level measured by the CECE system can be calculated as:

$$\frac{\tilde{T}}{T} = \sqrt{\frac{2}{B_{IF}} \int_{f_1}^{f_2} \frac{Re\{\gamma_c(f) - \gamma_{bg}\}}{1 - Re\{\gamma_c(f) - \gamma_{bg}\}} df} \quad (4.1)$$

where B_{IF} is the intermediate frequency bandwidth (the bandwidth of the band-pass filters used in the CECE system), f_1 and f_2 define the frequency range over which to integrate the coherence (the frequency width of the turbulent feature), γ_c is the complex coherence function, and γ_{bg} is the background coherence.

The key component of this equation, the complex coherence function, γ_c , is defined as [108, 110, 111, 112, 113, 114, 115, 116, 117]:

$$\gamma_c(f) = \frac{G_{xy}(f)}{\sqrt{G_{xx}(f)G_{yy}(f)}} \quad (4.2)$$

where G_{xy} is the one-sided cross-spectral density function between channels x and y , and G_{xx} and G_{yy} are the autospectral density functions of channels x and y , defined as:

$$G_{ij}(f) = 2F_i^*(f)F_j(f) \quad (4.3)$$

where F_i is the Fourier transform (frequency spectrum) of channel i calculated with a 50% overlapping Hanning window, ensemble averaged fast Fourier transform (FFT).

The background coherence, γ_{bg} , is calculated as:

$$\gamma_{bg} = Mean[\gamma_c]_{f_3}^{f_4} \quad (4.4)$$

where f_3 and f_4 define a frequency range far above the turbulent signal.

As is apparent from the above equations, one strength of this method of calculating the total temperature fluctuation level is that no absolute calibration of the radiometer is necessary. Since everything is normalized in Equation 4.1, only relative changes in signal matter, not the absolute value.

As noted above, these final forms of the equations used for the CECE analysis are not the only possible options. Early on in the work presented in this thesis, a less general form of Equation 4.1 was used:

$$\frac{\tilde{T}}{T} = \sqrt{\frac{2}{B_{IF}} \int_{f_1}^{f_2} \gamma_{xy}(f) df} \quad (4.5)$$

where all quantities are defined as above, and γ_{xy} is the magnitude of the complex coherence function:

$$\gamma_{xy}(f) = |\gamma_c(f)| = \sqrt{\frac{|G_{xy}(f)|^2}{G_{xx}(f)G_{yy}(f)}} \quad (4.6)$$

This form of the equation for the total temperature fluctuation level was used in many previous CECE references [71, 75, 76, 80, 105], as well as the the L- and I-mode analysis from Reference [100] that was performed as part of this thesis. The final version of the L- and I-mode analysis presented in this thesis in Chapter 5, however, has been updated with the more general formulas derived here, though the new analysis does not significantly change any of the results.

There are three key differences between Equation 4.1 and Equation 4.5.

1. The use of γ_c as opposed to γ_{xy} .
2. The use of $\gamma_c/(1 - \gamma_c)$ as opposed to just γ_c in the integrand.
3. The subtraction of the background, γ_{bg} .

As is presented in the derivation in Appendix C, all three of these differences make Equation 4.1 more general and more rigorous than Equation 4.5. Specifically,

using the real part of γ_c as opposed to the magnitude (aspect 1) is generally more mathematically rigorous, as described in Reference [115]. Using $\gamma_c/(1 - \gamma_c)$ instead of just γ_c (aspect 2) results from keeping an extra term in the perturbation expansion used to derive Equation 4.1. This essentially relaxes the assumption that the temperature fluctuation is much smaller than the thermal noise in the system. Finally, subtracting the background coherence (aspect 3) results from allowing for some common background noise between the two CECE channels. Such noise could be caused, for example, by finite filter overlap between the two CECE channels.

In addition to the derivation presented in Appendix C, Section 4.2.2 will use both synthetic and real experimental data to illustrate the advantages of using Equation 4.1. Interestingly, it appears that the more general formulation, especially background subtraction, significantly improved the results from ASDEX Upgrade, while it had little effect on the results from Alcator C-Mod. Possible reasons for this discrepancy will be given in the relevant sections below.

While this completes the equations required for the calculation of the temperature fluctuation level measured by the CECE diagnostic, for the purposes of validation it is vital to be able to quantify the uncertainty in this measurement. For this reason care was taken to accurately quantify the uncertainty in both γ_c , in order to compare spectra to transport models, and \tilde{T}/T , in order to compare the total fluctuation level.

To begin, Reference [109] has derived the standard deviation of the complex coherence function when calculated with ensemble averaging windows (as is the case here):

$$\sigma_{\gamma_c}(f) = \sqrt{\frac{1}{2n_d}(1 - |\gamma_c(f)|^2)^2} \quad (4.7)$$

where n_d is the number of independent ensemble averaging windows.

Previous literature discusses the sensitivity limit of the complex coherence function, which is typically defined as the level below which the signal is considered to

be indistinguishable from noise. This work will define this level as two standard deviations above zero.

One can then propagate this uncertainty through Equation 4.1 in order to obtain the approximate uncertainty in \tilde{T}/T . This derivation, given in Appendix C.3, assumes that γ_c is small compared to 1 (higher order terms are ignored), and gives:

$$\sigma_{\tilde{T}/T} \approx \frac{1}{\tilde{T}/T} \frac{1}{B_{IF}} \sqrt{\sum_i (\sigma_{\gamma_i} \delta f)^2} \quad (4.8)$$

where δf is the frequency resolution of the Fourier transform. This equation is expressed as a finite sum over frequency bin i , instead of an integral, since in practice this is how the uncertainty is generally calculated.

Finally, one would also like to calculate the sensitivity limit of the integrated fluctuation level, or, in other words, the smallest measurable total temperature fluctuation level. This is done by appropriately integrating the standard deviation of the complex coherence function from Equation 4.7, over the signal bandwidth. For the purposes of this calculation, the standard deviation of a signal with no coherence is used, as this corresponds to integrating the effects of pure noise. This integration, given in Appendix C.4, results in:

$$\frac{\tilde{T}}{T} \Big|_{limit}^{stat} > \sqrt{\frac{2}{\sqrt{N}} \frac{B_{sig}}{B_{IF}} \sqrt{\frac{B_{samp}}{2 \cdot B_{sig}}}} \quad (4.9)$$

where N is the total number of data points, $B_{sig} = f_2 - f_1$ is the signal bandwidth, B_{samp} is the total sampling rate, and B_{IF} is again the intermediate frequency bandwidth. For 1 second of data collection on ASDEX Upgrade, typical values would be: $N = 4 \times 10^6$, $B_{sig} = 10^5$, $B_{IF} = 10^8$, and $B_{samp} = 4 \times 10^6$. This gives $\tilde{T}/T|_{limit}^{stat} = 0.2\%$.

This limit differs slightly than that given in some previous literature [71, 120], but when B_{sig} is increased to the Nyquist Frequency, $B_{sig} = (1/2)B_{samp}$, Equation 4.9

reduces to the ideal radiometer limit given in References [71, 120]:

$$\frac{\tilde{T}}{T} \Big|_{limit}^{rad} > \sqrt{\frac{2}{\sqrt{N}} \frac{B_{sig}}{B_{IF}}} \quad (4.10)$$

This difference, originating from the derivation of Equation 4.9, is consistent with suggestions in References [73, 121] that the ideal radiometer limit must be corrected when the sampling rate is significantly higher than the signal bandwidth.

Just as with the equations used to calculate \tilde{T}/T , earlier results during work presented in this thesis did not utilize the most rigorous versions of these formulas. In particular, the work concerning L- and I-mode presented in Reference [100] utilized Equation 4.10, and not Equation 4.9, as the instrumental sensitivity limit, as the more rigorous derivation had not yet been completed. The revised results presented in this thesis in Chapter 5, however, do use the updated analysis, though the differences are only minor.

In addition, the work concerning L- and I-mode plasmas [100] also included in its calculations a statistical sensitivity limit, given by the bias error in the coherence calculation [115]:

$$\frac{\tilde{T}}{T} \Big|_{limit}^{bias} > \sqrt{\frac{B_{sig}}{B_{IF}} \int_{f_1}^{f_2} b(f) df} \quad (4.11)$$

where:

$$b(f) = \sqrt{\frac{1}{n_d} (1 - \gamma_{xy}^2)^2} \quad (4.12)$$

in which n_d is again the number of independent ensemble averaging windows.

In some instances this bias error limit was actually greater than the radiometer limit. In these cases, Equation 4.11 was cited as the lowest measurable total fluctuation level, as opposed to Equation 4.10. It turns out, however, that using the improved formula for \tilde{T}/T , Equation 4.1, as opposed to the old Equation 4.5, makes

this bias error limit irrelevant. For this reason, all calculations made with Equation 4.1 use only the more rigorous sensitivity limit given by Equation 4.9, and do not also calculate a bias error-based sensitivity limit.

4.2.2 Testing of Data Analysis Techniques

In addition to the derivations given in Appendix C, the equations given in the previous section were thoroughly tested using both synthetic and real experimental data [81].

First, synthetic data testing shows the advantages of using Equation 4.1 over Equation 4.5. Specifically, this testing addresses the use of γ_c as opposed to γ_{xy} and the use of $\gamma_c/(1 - \gamma_c)$ as opposed to just γ_c in the integrand (aspects 1 and 2 from the previous section). Since the data generated was entirely synthetic, the background coherence, γ_{bg} , was set to zero.

For the purposes of this testing, the synthetic data consisted of two, zero mean, random gaussian signals, both of which contained 1000 bins of 1024 points, equivalent to roughly one second of data collected at 1MHz (assuming that there was no overlap in the data bins). A common broadband feature, generated by low-pass filtering a random gaussian signal at a given frequency, was adjusted to be between 0.1% and 50% of the total signal amplitude and was added to both signals. The bandwidth of this common broadband feature was adjusted to be between 20 and 300 kHz for different tests, as will be shown later.

This set of synthetic data simulates two CECE channels, both of which contain a small correlated temperature fluctuation on top of uncorrelated noise. Both the bandwidth of the common signal and the amplitude were scanned in order to test the effects of amplitude and bandwidth on the output of the various calculation methods. In particular, the scan was intended to examine possible issues with integrating only γ_{xy} at both very large and very small bandwidths.

This synthetic data was then analyzed using variations of Equations 4.1 and 4.5,

each with a different integrand. Assuming that there was no coherent background ($\gamma_{bg} = 0$), the four integrands that were used were: γ_{xy} , $\gamma_{xy}/(1 - \gamma_{xy})$, $Re\{\gamma_c\}$, and $Re\{\gamma_c\}/(1 - Re\{\gamma_c\})$. This testing therefore isolates the effects of using γ_c or γ_{xy} and using $Re\{\gamma_c\}$ or $Re\{\gamma_c\}/(1 - Re\{\gamma_c\})$.




Figure 4-5: Calculated values of the total fluctuation level of the synthetic data, obtained using various integrands within Equation 4.1. The synthetic data here has a true fluctuation level of 1%, represented by the horizontal black dashed line. Integrals of γ_{xy} (blue), $\gamma_{xy}/(1 - \gamma_{xy})$ (cyan), $Re\{\gamma_c\}$ (green), and $Re\{\gamma_c\}/(1 - Re\{\gamma_c\})$ (red) are tested. $Re\{\gamma_c\}/(1 - Re\{\gamma_c\})$ most accurately recovers the true fluctuation level over the widest range of common signal bandwidths. Figure from [81].

The integration bounds used in Equation 4.1, f_1 and f_2 , were adjusted along with the bandwidth of the common broadband feature that was added to both signals such that the integral captured the entire signal without integrating significantly higher frequencies. These calculations used a synthetic B_{IF} of 200 MHz, though a choice of 100 MHz does not alter the conclusions drawn.

Figure 4-5 shows the results of the synthetic data analysis for an input common signal fluctuation level of 1%, plotting the calculated fluctuation level, \tilde{T}/T against the common signal bandwidth for the four different integrands. Ideally, the calculation method would recover the input fluctuation level of 1%, which is shown as the horizontal black dotted line, at all bandwidths. The result of each calculation method

is shown as a different color and pattern in the figure.

One can draw two primary conclusions from the results shown in Figure 4-5. First, one may note that at large signal bandwidths, utilizing either γ_{xy} or $\gamma_{xy}/(1-\gamma_{xy})$ tends to increasingly overestimate the fluctuation level as the signal bandwidth increases. This can be understood as follows. Since γ_{xy} is a magnitude, it is a positive definite quantity. Therefore any signal with finite noise will always have finite γ_{xy} even when there is no coherent fluctuation between the two channels. This finite level of γ_{xy} due to noise will be integrated along with the actual common signal. At small bandwidths, the common fluctuation will dominate, as it will be much larger than the noise level of γ_{xy} . As the bandwidth increases, however, the noise will become increasingly important, since its relative amplitude will be larger as the common fluctuation is spread out over a larger frequency range. This applies equally to γ_{xy} and $\gamma_{xy}/(1-\gamma_{xy})$. On the other hand, since $Re\{\gamma_c\}$ properly accounts for the phase between the two signals and can take negative values, it will not accumulate noise as the bandwidth is increased.

Second, the results shown in Figure 4-5 reveal that utilizing either γ_{xy} or $Re\{\gamma_c\}$ tends to underestimate the total fluctuation at small signal bandwidths, with the underestimation becoming more pronounced as the bandwidth continues to decrease. This effect originates from a different cause than the first. As is noted in Appendix C, a derivation of the formula for \tilde{T}/T that has only a numerator in the integrand makes the assumption that the auto-spectral density function of each input signal, $G_{ii}(f)$, is dominated by the background thermal noise, not by the temperature fluctuation signal. When the temperature fluctuation signal is spread out over a large bandwidth, this assumption holds true and using only $Re\{\gamma_c\}$ is reasonably accurate. When the signal bandwidth is rather small, however, the temperature fluctuation signal becomes significant in not only the cross-spectral density, but also in the auto-spectral density. In this case the normalization of the cross-spectral density by the auto-

spectral densities, which is the definition of both γ_c and γ_{xy} , tends to underestimate the total fluctuation level, since the auto-spectral densities are artificially raised above the background thermal noise.

This effect can be corrected for, however, by including $1 - Re\{\gamma_c\}$ in the denominator of the integrand. More rigorously, as shown in Appendix C, this denominator arises from relaxing the assumption that thermal noise dominates the auto-spectral density. Practically, this change is valid for any pair of signals that has nonzero noise, and therefore has a coherence less than 1.0. Perhaps unsurprisingly, all real experimental measurements have some finite level of noise, and therefore have a coherence less than 1.0.

While Figure 4-5 shows results from synthetic data with a common broadband fluctuation of 1%, as mentioned above, fluctuation levels between 0.1% and 50% of the total signal amplitude were tested, in addition to data with no common broadband signal. All testing with finite signal level showed similar results to those shown in Figure 4-5, with higher signal levels only further favoring $Re\{\gamma_c\}/(1 - Re\{\gamma_c\})$ as the most accurate calculation method. More specifically, a larger fluctuation level significantly exacerbates the second issue mentioned above (concerning the effect of the fluctuation on the auto-spectral density), but does reduce the effects of first problem (concerning integration of noise). The overall effect, however, is strongly deleterious, in that the reduction of the first issue only mildly offsets the exacerbation of the second issue.

On the other hand, analyzing synthetic data with zero common signal using γ_{xy} and $\gamma_{xy}/(1 - \gamma_{xy})$ still shows finite fluctuation levels due to integration of noise, as the first issue listed above becomes more severe. Using $Re\{\gamma_c\}$ and $Re\{\gamma_c\}/(1 - Re\{\gamma_c\})$, however, gives results that scatter near zero, as one would expect.

The results of the synthetic data testing shown in this section therefore complement the more rigorous analytic derivation shown in Appendix C. The derivation in

Appendix C shows that using $Re\{\gamma_c\}/(1 - Re\{\gamma_c\})$ in the integrand of Equation 4.1 is more rigorous and generally valid than previous formulations. Similarly, the results of synthetic data testing presented here show that Equation 4.1 more effectively recovers a known fluctuation level than previous equations.

Before moving on to background subtraction, it is noted that as mentioned above, the new calculation method had a significantly larger impact on the results from ASDEX Upgrade than those from Alcator C-Mod. While the details of these results will be presented in Section 4.3 and Chapters 5 and 6, it is briefly mentioned that most of the Alcator C-Mod measurements contained turbulent signals with bandwidths of between 100 and 200 kHz, while the ASDEX Upgrade results exhibited a greater range of bandwidths. As is apparent from Figure 4-5, utilizing γ_{xy} as opposed to $Re\{\gamma_c\}/(1 - Re\{\gamma_c\})$ in this bandwidth range typically results in a discrepancy of less than 5%, which is comparable to the uncertainty in the calculation anyway. It seems, then, that the Alcator C-Mod results fortuitously fall in a bandwidth range for which the more general Equation 4.1 is roughly equivalent to Equation 4.5. The lack of a coherent background in the Alcator C-Mod data will be addressed below.

In addition to the synthetic data testing that was used to address differences 1 and 2 between Equations 4.1 and 4.5, testing using noise source input and real experimental data was used to address the third difference, concerning the subtraction of background coherence. Again, the results presented here compliment the analytic derivation presented in Appendix C.

As is stated in Appendix C, the inclusion of γ_{bg} in Equation 4.1 results from allowing the two channels to have some level of common noise that is retained after the two channels are correlated. This common noise can come from a variety of sources, including electronic noise and finite filter overlap, among others.

While in many past CECE systems this term has proven to be unimportant, such as the system on Alcator C-Mod, it was discovered early in the operation of the CECE

system on ASDEX Upgrade that some pairs of nearby channels exhibited a non-zero complex coherence function, even when there was not expected to be any plasma turbulence. In situations where there was a clear fluctuation signal, the complex coherence function would remain more than two standard deviations above zero even at frequencies far above the turbulent feature.

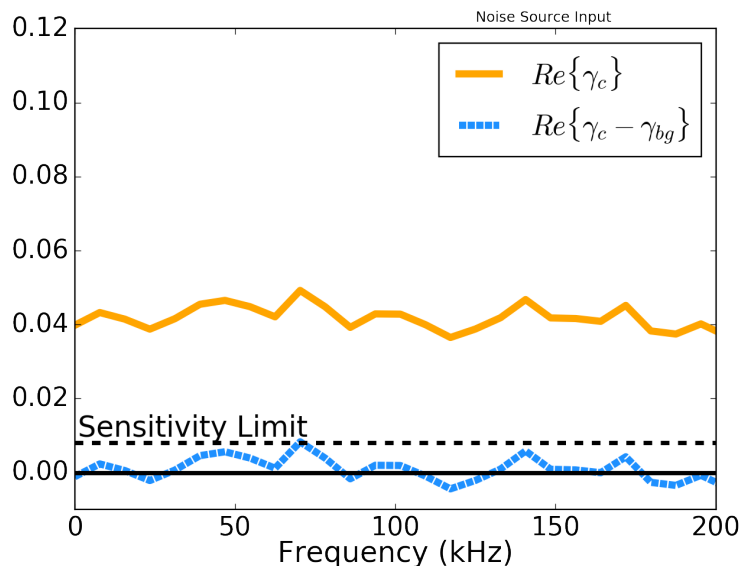


Figure 4-6: Noise source testing of the CECE system on ASDEX Upgrade. The orange solid line is before background subtraction and the blue dotted line is after. The largely flat nature of the background coherence suggests that finite filter overlap may be responsible for the noise. In addition, this flatness allows for effective subtraction of the same noise level at all frequencies.

One method of measuring this effect is to inject a noise source into the system, rather than measuring real plasma signal. It turns out that in the case of the system on ASDEX Upgrade, the noise generated by the RF amplifiers was sufficiently large to serve the purpose of a noise source input to the IF section. This meant that such noise source testing could be readily performed without removing the system from its rack or changing any of the cables. The results of such a test are shown in Figure 4-6 (the effect on experimental data will be shown later in Figure 4-7).

Note that the signal without background subtraction shown in Figure 4-6, which consists solely of correlated noise, is nearly flat across the frequency spectrum at

frequencies of interest (roughly 0 to 200 kHz). This has two main consequences. First, the broadband flat nature of the noise suggests that finite filter overlap may be a possible cause. Finite filter overlap means that the two bandpass filters in the IF section that separate the two measurement locations radially have some finite overlap in their passbands. This explanation is also consistent with the observation that the background level varies significantly depending on which pair of channels is observed, since such passband overlap depends on the quality of each individual filter's construction. In other words, even though the 3dB cutoff point of the filters are theoretically separated by 25 MHz, the drop-off of the passband may not be sufficiently sharp to avoid any overlap. This drop off also varies between filters, so some pairs would overlap more than others. See Reference [74] for further discussion of finite filter overlap in CECE systems.

Second, the flat nature of the coherent background noise means that a constant value can be subtracted from the complex coherence function at all frequencies. This greatly simplifies the analysis, and leads to the formulation presented in Equations 4.1 and 4.4. In the data presented in Figure 4-6, the background is averaged from 120 to 200 kHz, and is then subtracted as a constant from the complex coherence function at all frequencies.

After noise source testing, the impact of background subtraction on real experimental data was evaluated. Figure 4-7 shows the results of subtracting coherent background from experimental data on ASDEX Upgrade, from a pair of channels at $\rho_{tor} = 0.64$ in discharge 33995. As discussed above, Figure 4-7 (a) shows that without background subtraction the complex coherence function never drops to below the sensitivity limit, even at frequencies clearly above the extent of the turbulent feature. After subtraction, however, the turbulent feature remains clearly above the sensitivity limit while the complex coherence drops below this level at higher frequencies. In the data presented here, the γ_{bg} was calculated as the mean value of the complex

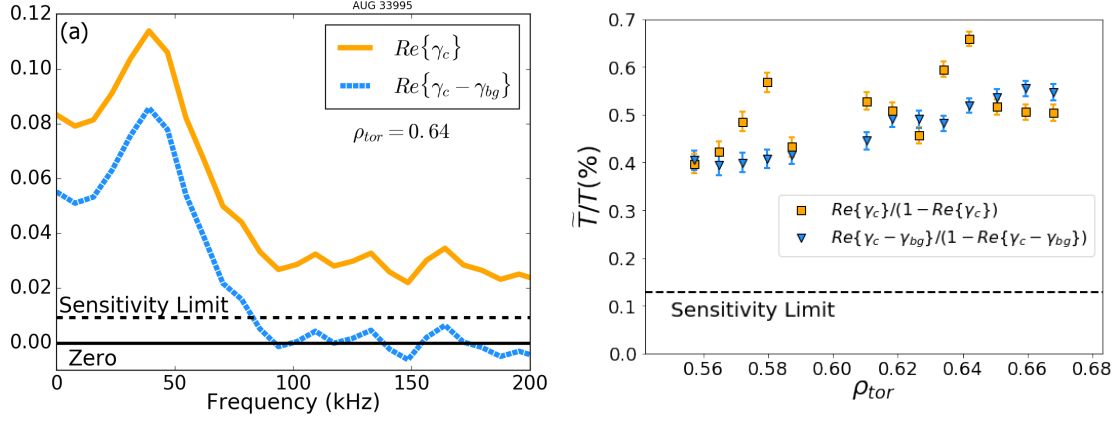


Figure 4-7: Effect of background subtraction on experimental data from ASDEX Upgrade. (a) shows the fluctuation spectrum and (b) shows the integrated fluctuation radial profile. Orange solid line in (a) and squares in (b) are the signal without background subtraction. Blue dotted line in (a) and triangles in (b) are the signal after background subtraction. Error bars in (b) represent one standard deviation. The black horizontal dotted line is the sensitivity limit. Figure from [81].

coherence function between 120 and 200 kHz.

Figure 4-7 (b) shows the integrated fluctuation radial profile from the same discharge, where each point is calculated using Equation 4.1 for a different pair of channels. The error bars in (b) represent one standard deviation, as calculated with Equation 4.8. The sensitivity limit, represented as the horizontal black dotted line, is calculated with Equation 4.9. Coherent background noise clearly has a significant impact on the radial fluctuation profile, creating large discrepancies in the total fluctuation level of channel pairs that are very close radially. The resulting structure in the radial profile appears to be clearly unphysical. On the other hand, subtracting the coherent background from each channel pair greatly smooths the overall profile. While one should not use only experimental data to justify a calculation technique, the resulting smooth profile is to some extent comforting, as this is what one might otherwise anticipate.

It was noted earlier that CECE observations on Alcator C-Mod did not contain the same coherent background as was seen on ASDEX Upgrade. Differences in the

CECE system hardware on the two machines is consistent with the supposition that the origin of the background is finite filter overlap. As describe in Section 4.1.1, the 100 MHz bandwidth fixed filters on Alcator C-Mod had a center-to-center separation of 140 MHz, while the same filters on ASDEX Upgrade had a separation of 125 MHz. The tighter frequency separation was chosen on ASDEX Upgrade after benchtop testing suggested that such a separation would be sufficient, but the observation of the coherent noise on the digitized signal indicates that this assessment was likely inaccurate. Moving the filters closer by 15 MHz may very well have led to just enough filter overlap to be detectable.

This section has thus shown the benefits of using Equation 4.1 over the previously used 4.5. While Appendix C shows the origin of the differences between the two equations analytically, this section has shown that the new equation performs better on synthetic data, noise source data, and real experimental data. This section has also described why the new equation had a significant impact on the analysis of data from ASDEX Upgrade, while having a minimal impact on data from Alcator C-Mod.

4.3 Experimental CECE Measurements

This section presents a few examples of CECE data from both Alcator C-Mod and ASDEX Upgrade, focusing on purely experimental observations. Chapters 5 and 6 will later use CECE data as a validation metric for turbulent transport models. For each machine, a few general examples of data are shown in order to familiarize the reader with what the final results of CECE data analysis look like and how to interpret these results. In addition, CECE measurements were used to investigate the transition from L- to I-mode on Alcator C-Mod, showing that electron temperature fluctuation levels generally decrease after the transition from L-mode to I-mode. A possible exception to this trend is observed in particularly low current plasmas. The large number of channels on the ASDEX Upgrade system enables wide radial profile

measurements, in which the fluctuation level generally increases with increasing minor radius. Finally, the ASDEX Upgrade system was used to investigate the impact of resonant magnetic perturbations (RMPs) on the electron temperature fluctuation level in H-mode plasmas.

4.3.1 Alcator C-Mod

Since the hardware on both machines evolved over time, it is important to be clear on the timeline. The rebuilt CECE section on Alcator C-Mod was installed in September of 2015 (replacing the system described in References [75] and [76]). As described above, the new system installed on Alcator C-Mod consisted of only four channels (two fixed frequency and two tunable), set up in two pairs of two. This means that the fluctuation spectrum could only be measured at two radial locations at any given time. For this reason, generally only the spectra, and not a radial profile, are plotted.

Example Data and Interpretation

Figure 4-8 shows data from one Alcator C-Mod discharge. This particular discharge, 1160504023, was a low current, high power L-mode discharge with the following plasma parameters: 5.4 T toroidal field, 0.55 MA plasma current, 1.6 MW ICRH, and $1.1 \times 10^{20} \text{1/m}^3$ line averaged density. The spectra shown in the figure are averaged over 0.7s of plasma operation. These spectra are plotted as γ_c as defined in Equation 4.2 and γ_{bg} as defined in Equation 4.4. Uncertainty is represented by the shaded regions, and is calculated with Equation 4.7. The dotted black line shows two standard deviations above zero.

The two pairs of channels measured fluctuations at ρ_{tor} of 0.73 and 0.81. The total fluctuation levels and uncertainties, calculated using Equations 4.1 and 4.8 were $\tilde{T}/T = 0.49\% \pm 0.03\%$ and $\tilde{T}/T = 0.79\% \pm 0.02\%$, respectively. The sensitivity limit of these fluctuation levels, as calculated with Equation 4.9, was 0.26%.

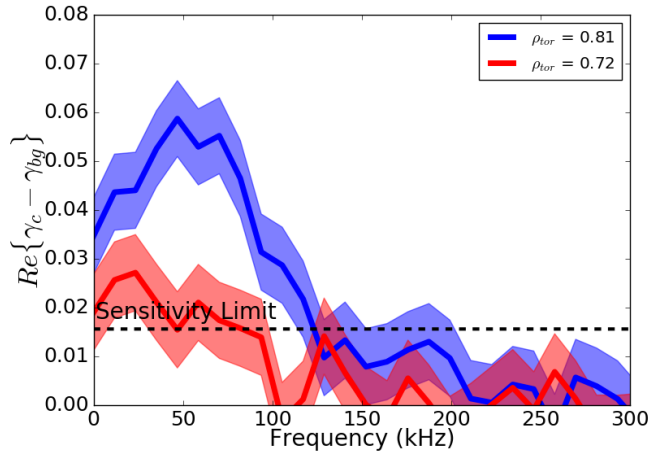


Figure 4-8: Electron temperature fluctuation spectra measured with the new CECE system on Alcator C-Mod in discharge 1160504023. Spectra are plotted as $Re\{\gamma_c - \gamma_{bg}\}$. Uncertainty is represented by the shaded regions. The horizontal dotted line represents two standard deviations above zero. Data is integrated from 0.7s of plasma operation.

One can therefore interpret the data shown in Figure 4-8 as follows. There are measurable fluctuations at both ρ_{tor} of 0.73 and 0.81, which surpass the integrated fluctuation level sensitivity limit by a fair margin. In terms of the spectra, the fluctuations at ρ_{tor} of 0.81 range in frequency from approximately 0 to 200 kHz, and remain above two standard deviations from zero until approximately 125 kHz. At ρ_{tor} of 0.73, the fluctuations are smaller in magnitude, and remain more than two standard deviations above zero only until approximately 100 kHz. The total fluctuation level increases at larger minor radius, as has been seen with many previous CECE measurements, and as is shown below for ASDEX Upgrade. Fluctuation spectra and integrated fluctuation levels such as those calculated here will be utilized later in this thesis in order to both explore experimental trends and to validate transport simulations.

L- and I-mode Plasmas

One of the earlier applications of the new CECE system on Alcator C-Mod was to investigate the difference in temperature fluctuation levels between L- and I-mode plasmas. First, some data from the old CECE hardware on Alcator C-Mod will be analyzed with the new analysis techniques described above in order to show the general trend of decreasing fluctuation after the transition from L-mode to I-mode. Note that this same discharge is one of those used in a validation study in Chapter 5. Next, data taken using the new hardware during a series of discharges will reveal a possible exception to this trend for particularly low-current (and thus poor confinement) I-mode discharges.

Consider first Alcator C-Mod discharge 1120921008, which is the discharge that is used in the validation of the gyrokinetic code GYRO in Chapter 5. The data presented here was originally published in Reference [100]. This plasma had the following parameters: $n_e = 0.7 - 0.8 \times 10^{20} \text{m}^{-3}$, $I_p = 1.1 \text{ MA}$, $B_t = 5.4 \text{ T}$, and $q_{95} = 3.2$. The $B \times \nabla B$ drift was away from the active X-point (generally known as unfavorable ∇B drift), enabling more robust access to I-mode [90]. In addition to ohmic heating, this plasma was heated with 1.6 MW of ion cyclotron range of frequency (ICRF) during the L-mode time period and 3.5 MW during the I-mode time period (the increase in power initiated the transition to I-mode). This is moderate power for Alcator C-Mod. This discharge was designed to have a clean transition from L-mode to I-mode in the same plasma in order to enable a very clear comparison of measurements and simulations. More information on the time traces and temperature and density profiles for this discharge will be given in Chapter 5.

Since this data was taken in 2012, the CECE hardware was still that described in References [75] and [76], and thus had six channels (three pairs). In order to be consistent when comparing the two operational regimes, 150 ms of time averaging was used for the CECE analysis in both the L-mode and I-mode portions of the discharge.




Figure 4-9: Turbulent electron temperature fluctuation spectra in L-mode (a) and I-mode (b) measured with the CECE system on Alcator C-Mod. The blue curves are at radial location $\rho_{tor} = 0.80$, green is at $\rho_{tor} = 0.84$, and red is at $\rho_{tor} = 0.89$. Uncertainty is represented by the shaded region around each line and the horizontal dotted line represents two standard deviations above zero. Data is taken from 150 ms of plasma operation at 1.1 MA plasma current. A similar figure, made with older data analysis techniques, is found in Reference [100].

In particular, the L-mode portion of the discharge that was analyzed was 0.85 to 1.00 s and the I-mode was 1.15 to 1.30 s. These time ranges were chosen such they were as close in time as possible while avoiding any transients from the L-I transition.

During the transition from L- to I-mode, the magnetic field, plasma current, and q_{95} were steady. The density increased slightly due to the increase in RF power, but a small change in density such as this is not uncommon in I-mode (primarily due to the increase in RF power required for the transition), and is entirely distinct from the significant rise in density associated with the transition to H-mode.

The fluctuation spectra from three radial locations in the L- and I-mode time periods are shown in Figure 4-9. Note that this is the same discharge and time periods as published in Reference [100], but the figure here is slightly different than Figure 3 in Reference [100] since the figure here utilizes the new calculation techniques developed in Reference [81] and described above, while Reference [100] utilizes older analysis techniques. These differences, however, are minor, and do not change any of

Location	L-mode $\frac{\tilde{T}}{T}$	I-mode $\frac{\tilde{T}}{T}$
$\rho_{\text{tor}} = 0.80$	$0.81 \pm 0.05\%$	$< 0.40\%$
$\rho_{\text{tor}} = 0.84$	$1.14 \pm 0.03\%$	$0.57 \pm 0.07\%$
$\rho_{\text{tor}} = 0.89$	$1.46 \pm 0.04\%$	$0.69 \pm 0.06\%$

Table 4.1: Integrated electron temperature fluctuation levels measured with CECE in L- and I-mode plasmas and calculated with Equation 4.1. The sensitivity limit is 0.40%. The background subtracted complex coherence function is integrated from 0 to 200 kHz in L-mode (with background calculated from 200 to 300 kHz) and from 300 to 500 kHz in I-mode (with background calculated from 650 to 750 kHz). This table shows similar results to that found in Reference [100].

the conclusions drawn in Reference [100].

The total fluctuation levels corresponding to the three radial locations shown in Figure 4-9 ($\rho_{\text{tor}} = 0.80, 0.84,$ and 0.89) are given in Table 4.1. Again, these values differ slightly from those given in Reference [100] due to the use of updated analysis techniques, but do not alter any of the central conclusions. As was described in the previous section, the differences are most noticeable at fluctuation levels just above the sensitivity limit and very large fluctuation levels, and are the smallest at moderate fluctuation levels.

These calculations reveal that the fluctuations decrease at all three radial locations after the transition from L-mode to I-mode. One of the I-mode radii is below the sensitivity limit of 0.4%, which still represents a decrease from the L-mode measurement, which is above sensitivity. This decrease in turbulence after the transition from L-mode to I-mode at all radii is consistent with observations in previous work [96, 97], as well as with decrease in turbulence drive terms in this discharge (which will be discussed in further detail in Chapter 5). These spectra will be compared to the outputs of gyrokinetic simulations in Chapter 5.

While in this discharge and in previous work electron temperature fluctuation

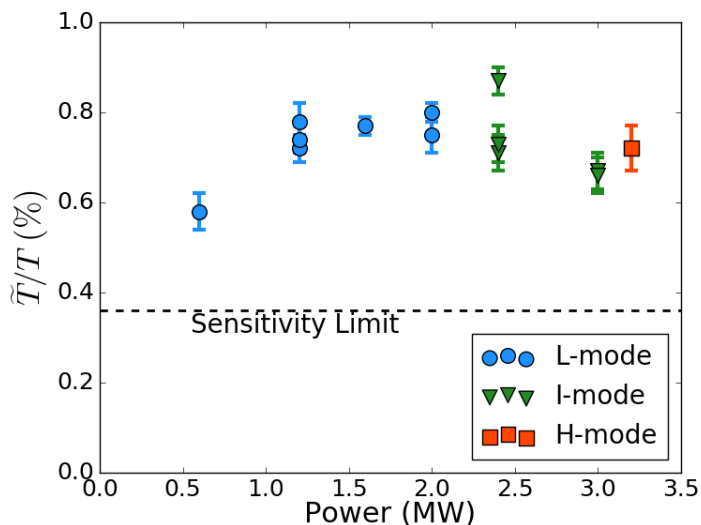


Figure 4-10: Total temperature fluctuation level, as calculated with Equation 4.1, plotted against ion cyclotron heating power for Alcator C-Mod L-, I-, and H-mode plasmas. L-mode plasmas are represented by the blue circles, I-mode by green triangles, and H-mode by red squares. Plasma current is 550 kA. Uncertainty is given by the error bars.

levels were generally observed to decrease after the transition from L-mode to I-mode, a few examples in which this is not the case have been observed. In particular, a series of very low current ($I_p \approx 550$ kA) L-I transition discharges were performed on Alcator C-Mod in order to investigate the effect of low current on temperature fluctuations in L-, I-, and H-mode plasmas, the results of which are shown in Figure 4-10. All of these discharges had an on-axis magnetic field of $B_t = 5.4$ T, line averaged density of $n_e \approx 1.1 \times 10^{20} \text{m}^{-3}$, and fluctuations were measured in the range $\rho_{tor} = 0.80 - 0.84$. ICRH power was scanned in order to effect the transition from L- to I- to H-mode. The data from Figure 4-10 is taken from five discharges (1160504019 to 1160504023), each with several time periods with different powers. CECE data was integrated for between 0.3 and 0.7 s for each time period. Uncertainty in the total temperature fluctuation level is represented by the error bars on each point.

Figure 4-10 shows a number of interesting trends. As is expected, as one increases the heating power, one eventually transitions from L-mode to I-mode. If the heating

power is further increased, the plasma then transitions to H-mode. One sees that as the power is increased in L-mode, the fluctuation level also increases. Interestingly, however, there is no sharp drop in temperature fluctuation level after either the transition to I-mode or the transition to H-mode, which is what has typically been observed in the past. Instead, just after the transition to I-mode, the fluctuation level is nearly identical to what it was in the higher power portion of the L-mode. Further increasing the power then slightly decreases the fluctuation level from its peak value, even after the transition to H-mode. Note that even though the relative fluctuation level (\tilde{T}/T) remains relatively constant after the transition to I-mode, the absolute fluctuation level (\tilde{T}) increases, as the background plasma temperature also increases after the transition.

One might speculate that this unusual behavior is related to the particularly low current in these discharges, since lowering the current is generally associated with reducing the confinement time and thus increasing heat transport in all three plasma regimes. To test this theory, the current in a repeat discharge was increased to 800 kA, and fluctuations were measured at heating powers of 1.6 and 2.0 MW, both of which were in the I-mode at the increased current. Indeed, at this higher current, the fluctuation level at both power levels was below the sensitivity limit of the CECE diagnostic of 0.34 %. This very low fluctuation level during a higher current I-mode is more in line with previously reported results.

While the observations presented here are interesting and may help shed light on the relationship between L- and I-mode plasmas, further investigation of this particular phenomenon is beyond the scope of this thesis, and will be the subject of future work. This concludes the experimental observations made with CECE on Alcator C-Mod. The diagnostic will later be used in Chapters 5 and 6 as a validation constraint for turbulent transport models.

4.3.2 ASDEX Upgrade

Consider now experimental observations from the CECE system on ASDEX Upgrade. The original ASDEX Upgrade hardware was installed in June of 2015, and the upgraded IF section was installed in segments between March and June of 2017 (first the amplifier-divider chassis and fixed frequency chassis were installed, and then the YIG chassis). All data taken since June 2017 utilized the expanded set of CECE channels described in Section 4.1.2.

Example Fluctuation Spectra and Radial Profiles

The CECE system on ASDEX Upgrade is capable of producing similar spectra to that on Alcator C-Mod, though the greatly increased number of channels and close spacing also enables it to produce finely spaced radial profiles of the total fluctuation level. To demonstrate this capability, data will be shown from a series of three repeat L-mode plasma discharges on ASDEX Upgrade (33995, 33996, and 33997) which operated at 2.5 T toroidal field, 1.0 MA plasma current, 1.0 MW ECH, and $2.0 \times 10^{19} \text{m}^{-3}$ line averaged density. This data was originally presented in Reference [81].

Figure 4-11 shows the analysis results from three seconds of plasma data in each of these ASDEX Upgrade discharges. As above, the fluctuation spectra shown in (a) are represented with γ_c and γ_{bg} from Equations 4.2 and 4.4. All other aspects of the figure are also the same as shown in Figure 4-8. While the system has 30 channels (and 16 of the same bandwidth with sequential center frequencies), spectra from only three pairs, located at $\rho_{tor} = 0.58$, 0.63, and 0.66, are shown in order to reduce clutter in the figure. When integrated from 20 to 100 kHz (known low frequency MHD activity is excluded in this manner) using Equation 4.1, these three pairs have total fluctuation levels of $\tilde{T}/T = 0.41\%$, 0.50%, and 0.57%, respectively. All three of these total fluctuation levels are far above the sensitivity limit of 0.13%, as calculated

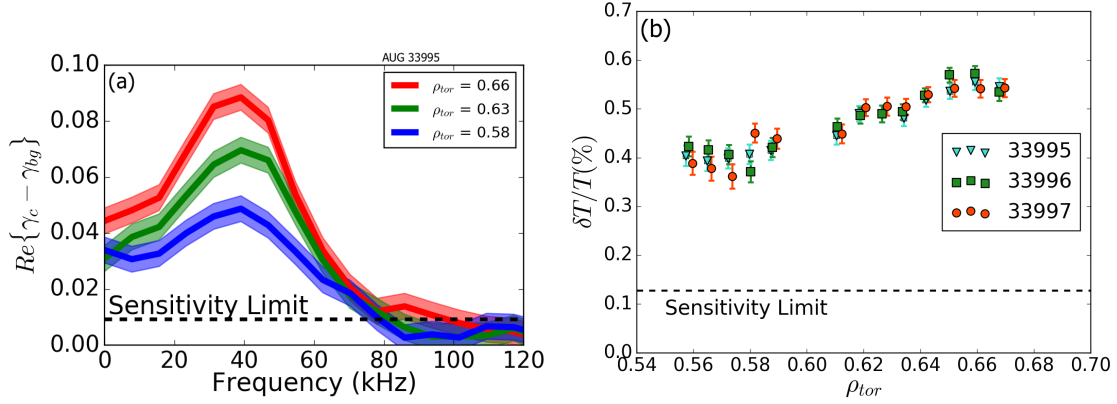


Figure 4-11: Electron temperature fluctuations measured with the new CECE system on ASDEX Upgrade. (a) fluctuation spectra, plotted as $Re\{\gamma_c - \gamma_{bg}\}$, at three radial locations in discharge 33995. Uncertainty is represented as the shaded regions. (b) radial profiles of the integrated fluctuation level from three repeat discharges, 33995, 33996, and 33997. Fluctuations were integrated from 20 to 100 kHz. Error bars represent one standard deviation of uncertainty. Data is integrated over 3s of plasma. Figure originally from [81].

with Equation 4.9.

When one expands this calculation of the total fluctuation level to include all of the channel pairs of the same bandwidth and spacing (except for one channel that was nonoperational), one can calculate the fluctuation level at 13 radial locations (16 channels gives 15 pairs, and one missing channel in the middle gives 13 pairs). The results of these calculations for all three ASDEX Upgrade discharges are shown in Figure 4-11 (b). Error bars in this plot represent one standard deviation, as calculated with Equation 4.8. The horizontal black dashed line is the sensitivity limit calculated with Equation 4.9.

Figure 4-11 reveals that in these three discharges, the total fluctuation level generally increases with increasing minor radius from $\rho_{tor} = 0.56$ to $\rho_{tor} = 0.68$. This observation is consistent with past observations on other machines [100, 76]. In addition, note that the profiles from the three repeat discharges generally agree within uncertainties. This result speaks to the robust repeatability of CECE measurements, engendering confidence that such measurements are accurately capturing plasma be-

havior from discharge to discharge.

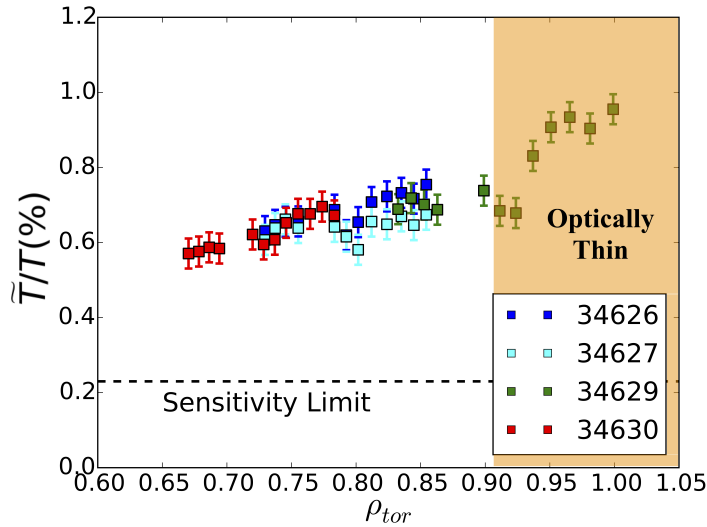


Figure 4-12: Temperature fluctuation profiles measured with the CECE system on ASDEX Upgrade during discharges 34626, 34627, 34629, and 34630. The four discharges were exact repeats except for slight tweaks to the magnetic field. The radial region outside of $\rho_{tor} \approx 0.9$ is optically thin, so ECE is not a reliable measure of temperature fluctuations.

Another example of the robustness of the measurements taken with the new ASDEX Upgrade CECE system is shown in Figure 4-12. This figure shows data taken from four repeat discharges, where only the magnetic field was varied by a few percent between discharges in order to scan the CECE measurement location. The resulting fluctuation profiles overlap radially, and show that the diagnostic can once again make very repeatable measurements between discharges. In addition, the varied magnetic field enables a larger radial range than is possible with a single discharge, revealing a consistent trend of increasing fluctuation level with radius. Note that the outermost radial points (outside of $\rho_{tor} \approx 0.9$, shaded in orange) are optically thin and are therefore not entirely reliable measures of the electron temperature fluctuation level.

H-mode Fluctuation Measurements

In addition to the L-mode measurements just shown, the CECE system on ASDEX Upgrade has also been able to measure fluctuations in H-mode plasmas, which has historically been very difficult due to the small fluctuation amplitudes. As such, H-mode fluctuation data has rarely been published previously, with Reference [78] as one of the only examples of such measurements. Note that the previous section also showed measurements made in one low current H-mode on Alcator C-Mod.

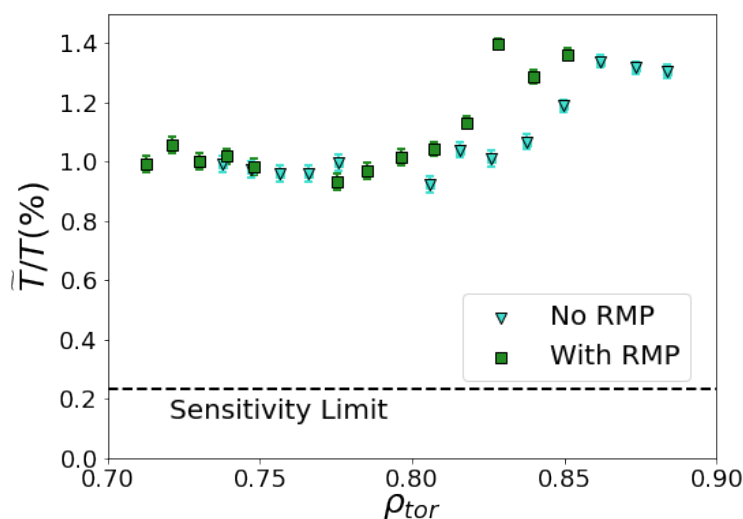


Figure 4-13: Temperature fluctuation level profiles measured with CECE during ASDEX Upgrade discharge 34577. Turquoise triangles are without RMPs and green squares are with RMPs.

In particular, measurements were made in an H-mode discharge (34577) during which resonant magnetic perturbations (RMPs) were applied. Fluctuations were measured both before and after the RMPs were applied. This discharge operated with the following parameters: $B_t = 2.6$ T, $I_p = 0.6$ MA, $n_e = 4.0 \times 10^{19} \text{m}^{-3}$, $P_{NBI} = 5.0$ MW, and $P_{ECH} = 1.7$ MW. CECE data was taken from 1.3 to 2.0 s without RMPs, and then from 2.0 to 2.7 s with RMPs. The radial fluctuation profiles from this discharge are shown in Figure 4-13.

It was somewhat surprising that fluctuations were measurable at all in this H-

mode plasma, as H-modes generally have very low fluctuation levels. One possible explanation, however, is that just like the I-modes and H-mode described above on Alcator C-Mod, this H-mode was relatively low current at only 600 kA. This low current leads to relatively poor confinement, even in an H-mode (the discharge has $H_{98} \approx 0.7$), which may be responsible for the increased fluctuation levels. In addition, these measurements show a clear difference between the time period with and without RMPs. In particular, Figure 4-13 shows that the fluctuation levels are very similar in both time periods, with the clear exception of the region between $\rho_{tor} \approx 0.82$ and $\rho_{tor} \approx 0.85$. In this region, which is relatively close to the top of the pedestal at $\rho_{tor} \approx 0.90$, the time period with RMP's has clearly increased fluctuation levels compared to the time period without.

Again, these measurements are primarily presented as an example of the capability of the CECE system on ASDEX Upgrade, but may be the subject of future investigation. Temperature fluctuations in H-mode are not nearly as well documented as those in L-mode, so this may be a rich area for further research.

Conclusions

This chapter has presented an overview of the CECE diagnostics on Alcator C-Mod and ASDEX Upgrade that were constructed as part of this thesis, and has given examples of the experimental data that has been collected on both machines. The chapter began with a brief overview of how heterodyne radiometers in general work, and how a correlation electron cyclotron emission diagnostic can be constructed as a radiometer. The specific hardware on Alcator C-Mod and ASDEX Upgrade was then described. The optics and RF section on Alcator C-Mod were constructed by a previous student, but the IF section was entirely new. Two different versions of the IF section for ASDEX Upgrade were installed as part of this thesis, an early version with six channels and an upgraded system with 30 channels.

In addition to the hardware on Alcator C-Mod and ASDEX Upgrade, new analysis techniques to extract the temperature fluctuation spectra and integrated temperature fluctuation levels were developed as part of this thesis. These new techniques are generalizations of previous analysis techniques, which relax several assumptions and therefore account more completely for various complications that may arise when analyzing CECE data. Both synthetic data and real experimental data were utilized to show the validity and superiority of these new analysis techniques. Finally, this chapter showed experimental data from both Alcator C-Mod and ASDEX Upgrade in order to demonstrate the capabilities of the two systems and the analysis technique, and to note some interesting experimental phenomena that may be the subject of future work. The CECE diagnostics and analysis techniques described in this chapter will be used in the validation of turbulent transport models in Chapters 5 and 6.

Chapter 5

Validation of Gyrokinetic Codes

This chapter describes validation studies of the gyrokinetic codes GYRO [18] and GENE [9] on Alcator C-Mod and ASDEX Upgrade, respectively. Chapter 2 motivated the importance of the validation of all turbulent transport models, and introduced the concept of using multiple validation constraints in this process. Chapters 3 and 4 described the development of ways to measure validation constraints, specifically the perturbative diffusivity and electron temperature fluctuations. This chapter and the next will now put those measurements to use, validating a variety of turbulent transport models. Of the four validation studies presented in this chapter, two were performed entirely as part of this thesis work, and two were performed outside of the thesis directly, but contained significant contributions from this thesis in terms of analysis and hardware.

This chapter will begin with two validation studies of GYRO on Alcator C-Mod. The first validation study looks at both L- and I-mode plasmas on Alcator C-Mod and uses both perturbative diffusivity and temperature fluctuations as constraints. The second study focuses on the addition of perturbative diffusivity as a validation constraint to several previously performed validation efforts, including studies with both ion-scale and multi-scale GYRO simulations.

The chapter will then present two validation studies of GENE on ASDEX Upgrade, focusing on two radial locations in the same discharge. The first, focusing on $\rho_{tor} = 0.49$, uses perturbative diffusivity (χ_e^{pert}) as a constraint. The second, focusing on a radial location of $\rho_{tor} = 0.75$, uses temperature fluctuations (\tilde{T}_e/T_e), the density-temperature fluctuation phase angle (α_{nt}), and electron temperature fluctuation radial correlation length (L_r) as constraints.

After presenting the results of all four validation efforts, this chapter will discuss overall implications and how these studies fit into a larger validation picture that has developed over the years. Specifically, the question of when multi-scale effects are important, and when ion-scale simulations are sufficient to model the experimental conditions, will be addressed. Although the results from this chapter are not themselves able to give an answer to this question, they do set the stage for the multi-machine, multi-discharge validation study with TGLF that is presented in Chapter 6, which was done explicitly in order to address the importance of multi-scale effects.

5.1 Validation of GYRO on Alcator C-Mod

Consider first the validation of the gyrokinetic code GYRO on Alcator C-Mod. This section will present two validation studies. The first study was performed entirely as part of this thesis, and focused on an L-mode time period and an I-mode time period within the same discharge. Global, ion-scale GYRO simulations were performed on these plasmas, using heat fluxes, perturbative diffusivity, and electron temperature fluctuation levels as validation constraints. This study shows that ion-scale gyrokinetic simulations robustly disagree with experimental measurements in this discharge, motivating investigation of when multi-scale effects are important, and if such effects could account for the discrepancies in this study.

The second study was not performed entirely as part of this thesis, but this thesis work contributed the perturbative diffusivity measurements as a validation constraint.

It focuses on several Alcator C-Mod discharges, adding the constraint of the perturbative diffusivity to other previously performed validation efforts, including one that used both ion- and multi-scale simulations. One of the main results of this validation study is that the perturbative thermal diffusivity is a powerful validation constraint that can differentiate ion- and multi-scale simulations.

5.1.1 L- and I-mode Validation

The first validation study in this chapter is fairly traditional, in the sense that it performs many experimental measurements in order to obtain validation constraints, and then compares these experimental constraints to the output of nonlinear ion-scale gyrokinetic simulations. The novel aspects of this study include the focus on I-mode plasmas and the use of four validation constraints: ion heat flux, electron heat flux, electron temperature fluctuation levels, and perturbative thermal diffusivity. This study will show that ion-scale simulations robustly disagree with both L- and I-mode plasmas when all validation constraints are applied.

First, a brief reiteration of why I-mode is an interesting target for a validation study. Recall from Chapter 3, that I-mode is a naturally ELM-free high-performance regime in which there is a temperature pedestal, but no density pedestal [89, 90, 91, 122]. The overall energy confinement time is often comparable to or longer than similar H-modes. In addition, I-mode does not suffer from core impurity accumulation, even in a metal walled machine, an issue that often plagues H-modes [90, 123, 124]. Though I-mode has also been observed on ASDEX Upgrade and DIII-D, Alcator C-Mod's high magnetic field made it uniquely capable of robustly accessing I-mode, as the power window for I-mode operation widens with increased magnetic field [91, 92, 93, 125].

Gyrokinetic simulations of I-mode plasmas are rare, but a few have been performed. One set of past studies showed that ion-scale GYRO simulations were able

to match the ion and electron heat fluxes within experimental uncertainty in an I-mode plasma at $\rho_{tor} = 0.8$, but not at $\rho_{tor} = 0.6$ [96, 97]. No additional validation constraints, however, were originally applied in these studies (some experimentally observed trends were compared, but not as rigorous validation constraints).¹

This section will now present a complete validation study performed specifically to investigate differences between L- and I-mode plasmas, both experimentally and with simulation, applying additional validation constraints. The work in this section was originally published in Reference [100].

The validation study here is based on global, nonlinear, ion-scale GYRO [18] simulations. This study makes use of all of the experimental methods and hardware described in both Chapters 3 and 4 in order to measure the experimental validation constraints. The study will show that in the plasma considered here, GYRO is able to match the experimental ion heat flux, but under-predicts the electron heat flux (at most radii), under-predicts electron temperature fluctuations, and under-predicts the perturbative diffusivity in both L- and I-mode. Linear addition of electron-scale simulation heat fluxes (which, to be fair, has been shown to be unrigorous in the past [12, 13, 14, 126]) was unable to resolve the discrepancy.

Experimental Measurements

The plasma under consideration for this validation study was Alcator C-Mod discharge 1120921008, which had the following plasma parameters: $\bar{n}_e = 0.7 - 0.8 \times 10^{20} \text{m}^{-3}$, $I_p = 1.1 \text{ MA}$, $B_t = 5.4 \text{ T}$, and $q_{95} = 3.2$. The discharge was run with the $B \times \nabla B$ drift away from the active X-point, which is generally referred to as the unfavorable (for H-mode) ∇B drift direction. The plasma was heated with 1.6 MW of ICRH in the L-mode time period and 3.5 MW of ICRH in the I-mode time period

¹Note that the next study presented in this thesis adds the perturbative diffusivity as a constraint to those past simulations, showing in Table 5.3 that the ion-scale simulations robustly under-predict the experimentally measured perturbative diffusivity, though the predicted qualitative trend that I-mode has a higher perturbative diffusivity than L-mode was captured [55].




Figure 5-1: The total RF heating power, central electron temperature, edge electron temperature, and line-averaged density for Alcator C-Mod discharge 1120921008. The L-mode period of this discharge that was analyzed is highlighted in blue. The I-mode period is highlighted in red. Figure originally from Reference [100].

(the increased heating power triggered the transition to I-mode).

Figure 5-1 shows time traces during the plasma discharge, revealing that the discharge had a clean transition from L-mode to I-mode that occurred between roughly 1.0 and 1.1 seconds. The discharge was specifically designed in order to have as clean of a transition from L- to I-mode as feasible, varying as few plasma parameters as possible. The RF ramp at around 1.0 seconds triggered the transition to I-mode leading to an increase in the temperature. There is also a slight increase in density associated with the increase in RF power, though this is entirely distinct from the sharp density rise associated with the transition to H-mode. Magnetic field, plasma current, q_{95} , and plasma shape were constant across this transition. All of the plasma parameters here are measured with the Alcator C-Mod diagnostics described in Chapter 2.

In order to avoid any complications involving the transition period, only steady periods of L-mode and I-mode were analyzed. As highlighted in Figure 5-1, for the




Figure 5-2: The time averaged electron temperature, ion temperature, and electron density profiles from the L- and I-mode portions of Alcator C-Mod discharge 1120921008, as well as their normalized gradient scale lengths. L-mode is blue and I-mode is red. Uncertainty is represented by the shaded regions. The green highlighted region indicates the region in which the gyrokinetic simulation was performed, and is the region for which the gradient scale lengths (a/L_{Te} , a/L_{Ti} , and a/L_{ne}) are shown. Figure originally from Reference [100].

remainder of this section the L-mode time period will refer to 0.85 to 1.0 s in this discharge, and the I-mode time period will refer to 1.15 to 1.3 s in this discharge. Note that for the rest of this validation study, in figures showing data from both time periods L-mode data will be blue and I-mode data will be red. Both of these periods had relatively stable background plasma parameters. The energy confinement time, τ_E , in the L-mode time period was approximately 30 ms ($H_{98,y2} = 0.65$), and in the I-mode time period it was approximately 40 ms ($H_{98,y2} = 0.95$). Both time periods therefore span multiple confinement times.

The electron temperature, ion temperature, and electron density profiles obtained by averaging over the L- and I-mode time periods are shown in Figure 5-2. This figure also shows the normalized gradient scale lengths for each of these parameters, $a/L_y = -(a/y)(dy/dr)$, where a is the plasma minor radius, r is radius, and y is

the parameter of interest. The normalized gradient scale lengths are drive terms for various turbulent modes, and are thus of particular interest when comparing L- and I-mode turbulent simulations.

This study will focus primarily on the region between $\rho_{tor} = 0.6$ and $\rho_{tor} = 0.9$ as this is well outside of the sawtooth mixing radius and inside of the edge-pedestal region. This region is highlighted in green in Figure 5-2 and is the region for which the normalized gradient scale lengths are shown. Inside of the sawtooth mixing radius, transport is heavily influenced by the sawtooth crash, and it is therefore generally unwise to apply a purely turbulent transport model. Applying gyrokinetic codes in the edge-pedestal region requires incredibly careful pedestal profile measurements, as well as great care in how the gyrokinetic model is used, and is therefore beyond the scope of this work. See Reference [123] for one example of edge-pedestal analysis in I-mode plasmas.

As displayed in Figure 5-2, the uncertainties in the normalized gradient scale lengths are 22% for a/L_{Te} , 40% for a/L_{Ti} , and 30% for a/L_{ne} over the analysis region, as calculated by combining diagnostic, fitting, and time uncertainty. These are the ranges over which the turbulence drive terms will be scanned for input to the GYRO simulations in the next section.

As a brief aside, one can make some qualitative conclusions about what the profiles in Figure 5-2 imply about differences in turbulence in the L- and I-mode plasmas. Very generally, Figure 5-2 shows that all three normalized gradient scale lengths (a/L_{Te} , a/L_{Ti} and a/L_{ne}), which drive turbulence, decrease after the transition to I-mode. This would seem to suggest that turbulence levels should decrease in the core of the I-mode plasma compared to the L-mode (especially considering the fact that the $E \times B$ shearing increases, as is shown later), a conclusion which is confirmed with CECE measurements.

Electron temperature fluctuations were measured with the Alcator C-Mod CECE

Location	L-mode $\frac{\tilde{T}}{T}$	I-mode $\frac{\tilde{T}}{T}$
$\rho_{\text{tor}} = \mathbf{0.80}$	$0.81 \pm 0.05\%$	$< 0.40\%$
$\rho_{\text{tor}} = \mathbf{0.84}$	$1.14 \pm 0.03\%$	$0.57 \pm 0.07\%$
$\rho_{\text{tor}} = \mathbf{0.89}$	$1.46 \pm 0.04\%$	$0.69 \pm 0.06\%$

Table 5.1: A copy of Table 4.1, reprinted here for convenience. Integrated electron temperature fluctuation levels measured with CECE in L- and I-mode plasmas and calculated with Equation 4.1. The sensitivity limit is 0.40%. The background subtracted complex coherence function is integrated from 0 to 200 kHz in L-mode (with background calculated from 200 to 300 kHz) and from 300 to 500 kHz in I-mode (with background calculated from 650 to 750 kHz). This table shows similar results to those found in Reference [100].

hardware and the analysis techniques described in Chapter 4. The frequency spectra measured in this discharge were shown earlier in Figure 4-9 as an example of the experimental capabilities of the Alcator C-Mod CECE system. The integrated fluctuation levels were calculated in Table 4.1, and are repeated here as Table 5.1 for convenience. Note again that Reference [100] used older analysis techniques, but the new analysis techniques developed as part of this thesis do not change any of the conclusions made in Reference [100]. The next section will directly compare the measured CECE crosspower spectrum to the output of gyrokinetic simulations, instead of the coherence, so the differences between the new and old analysis techniques are minor.

The perturbative diffusivity in this plasma was measured with the partial sawtooth heat pulse method as described in Chapter 3. In this particular case, the measurement is a radial average between $\rho_{\text{tor}} = 0.64$ and 0.84, chosen to avoid the region too close to the mixing radius and that so far out that the heat pulse amplitude is too small to measure. The perturbative diffusivity in L-mode was $\chi_e^{\text{pert}} = 4.0 \pm 0.6 \text{ m}^2/\text{s}$ and in I-mode was $\chi_e^{\text{pert}} = 4.3 \pm 0.9 \text{ m}^2/\text{s}$. While these values are the same to within the experimental uncertainty, the slight increase in perturbative diffusivity after the

transition to I-mode is consistent with the experimental results presented in Chapter 3 and Reference [55].

For completeness, the power balance diffusivities were also calculated (using TRANSP). The radially averaged power balance diffusivity in L-mode was $\chi_e^{PB} = 1.4 \pm 0.4 \text{ m}^2/\text{s}$ and in I-mode was $\chi_e^{PB} = 1.0 \pm 0.3 \text{ m}^2/\text{s}$. As is expected, the power balance diffusivity is lower in I-mode, indicating that it is a high confinement regime. Note that this is not quite as big as the traditional factor of two drop in power balance diffusivity seen after the transition from L- to H-mode [97]. This can easily be explained, however, by the fact that H-modes generally have a sharp density rise, and the power balance diffusivity depends inversely on the plasma density.

This completes the experimental measurements made in this discharge that will be used to validate the GYRO simulations.

GYRO Simulations, Linear Results, and Methodology

This section will describe the validation of global, nonlinear, ion-scale GYRO simulations of both the L- and I-mode portions of the discharge in question. This is largely the same methodology used in the validation of GENE on ASDEX Upgrade, and is a widely used workflow in the validation of gyrokinetic codes. As described above, the simulations are compared to heat fluxes, electron temperature fluctuations, and perturbative diffusivity.

These GYRO simulations were performed in the region $\rho_{tor} = 0.65 - 0.9$. They included three kinetic species (gyrokinetic ions and impurities, and drift-kinetic electrons), realistic geometry using the Miller parameterization, rotation effects (including $E \times B$ shear), and electron-ion and ion-ion collisions. Due to the low β of the discharge (and in particular the low gradient of β_{pol}), electromagnetic effects are negligible, so the simulations were electrostatic. These simulations were ion-scale, capturing long wavelength turbulence up to approximately $k_\theta \rho_s \approx 1.3$, where k_θ is the binormal

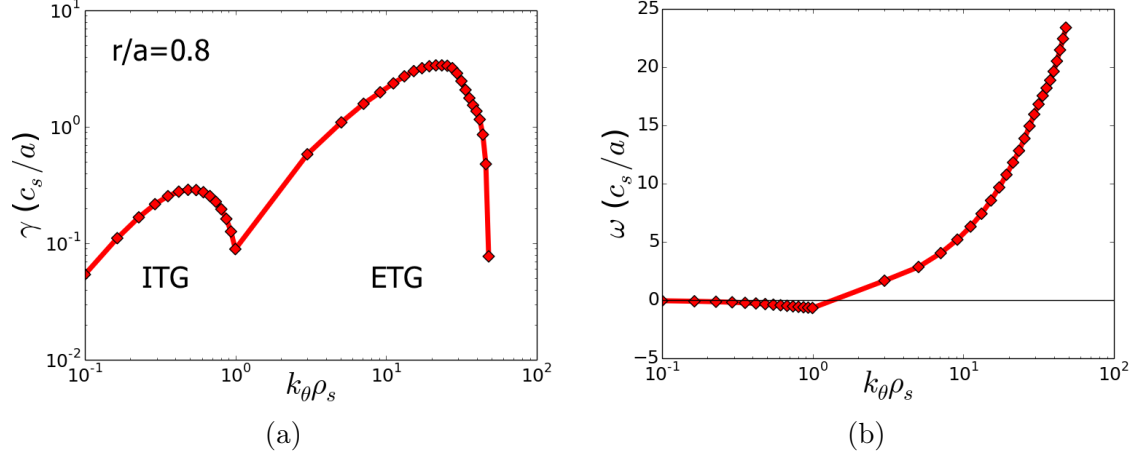


Figure 5-3: Linear growth rates (a) and real frequencies (b) at $\rho_{tor} = 0.8$ of the L-mode GYRO simulations. At long wavelengths, ITG dominates, while at short wavelengths ETG is dominantly unstable. The I-mode linear growth rate analysis shows similar results. Figure originally from Reference [100].

wavenumber and ρ_s is the ion Larmor radius evaluated at the electron temperature.

The simulation box size was approximately 105 by 120 ρ_s in the radial and binormal directions, and included 28 toroidal modes and approximately 500 radial grid points (for a grid spacing of $\sim 0.25\rho_s$). In order to perform global simulations, GYRO utilizes benign buffer regions. These buffers were approximately 11 (8) ρ_s wide on the inner (outer) region of for the L-mode simulation and 13 (8) ρ_s wide on the inner (outer) region of the I-mode simulation. The buffer region sources used default values, with the source annihilation rate, $\nu_{source} = 0.1c_s/a$. See Reference [95] for more information on buffer regions in the GYRO code.

A summary of the geometry and plasma parameter inputs to the GYRO simulations for both L- and I-mode time periods are given in Table 5.2. Local values are given at $\rho_{tor} = 0.8$, though the simulations themselves were global.

Consider first the linear stability of the plasmas in question. Linear growth rates and real frequencies for the L-mode plasma at $\rho_{tor} = 0.8$ are shown in Figure 5-3. Results for the I-mode are very similar. In GYRO, positive real frequencies are indicative of modes moving in the electron diamagnetic drift direction (electron modes),

Table 5.2: Input parameters for the GYRO simulations of the L- and I-mode plasmas described in this validation study. All values are the nominal experimental values, except for those in parentheses, which are for the ion heat flux-matched simulations. This table shows the values at $\rho_{tor} = 0.8$, though the values vary with radius throughout the global simulation domain. For the ion heat flux-matched simulations, a/L_{T_i} was adjusted down by 1% in L-mode and down by 29% in I-mode. Table originally from Reference [100].

	L-mode	I-mode
ρ_{tor}	0.8	0.8
$n_e(10^{20}m^{-3})$	0.94	1.14
T_e (keV)	0.76	1.26
a/L_n	1.67	0.97
a/L_{T_e}	5.81	3.60
a/L_{T_i}	4.79 (4.75)	3.18 (2.27)
$a/L_{T_{Imp}}$	4.79 (4.75)	3.18 (2.27)
T_i/T_e	1.12	0.83
n_D/n_e	0.758	0.830
n_{Imp}/n_e	0.018	0.015
Z_{eff}	3.90	2.70
$\nu_{ei} (a/c_s)$	0.29	0.13
$R_0(r)/a$	3.05	3.05
$\Delta=dR_0(r)/dr$	-0.08	-0.11
q	1.92	2.03
$\hat{s}=r \, d\ln(q)/dr$	2.13	2.13
κ	1.30	1.30
$s_\kappa=r \, d\ln(\kappa)/dr$	0.27	0.30
δ	0.20	0.21
$s_\delta=r \, d\delta/dr$	0.46	0.50
$\gamma_{E \times B} (a/c_s)$	0.0009	0.0524
$\gamma_p (a/c_s)$	0.006	0.405
$\rho^* = \rho_s/a$	0.0022	0.0028
$a/c_s (\mu s)$	0.91	1.17

while negative real frequencies are indicative of modes moving in the ion diamagnetic drift direction (ion modes). Note that the real frequencies in GYRO use the *opposite sign convention* as that used in GENE, which is the code used in Section 5.2.

At long wavelengths (low wavenumbers), the dominant mode has negative real frequency in GYRO, indicative of the ion temperature gradient (ITG) mode. At high wavenumber, the dominant mode has positive real frequency in GYRO, indicative of the electron temperature gradient (ETG) mode.

Since the ITG mode was dominantly unstable at low wavenumbers, this study followed the traditional gyrokinetic validation workflow and varied only a/L_{Ti} (within experimental uncertainty) in order to try to match the experimental heat fluxes. Such a method is consistent with past validation studies on Alcator C-Mod, in which ITG is typically the dominant mode at low wavenumbers (in L-, I- and H-modes) [14, 96, 97, 127, 128]. Test scans of the electron temperature gradient revealed that it had little impact on the heat fluxes, justifying such a method in this particular case. This is in contrast to the work presented later in Section 5.2, in which it was necessary to scan both the ion and electron temperature gradient in order to match either the ion or electron heat fluxes. The desire to scan both temperature gradients, as well as other inputs, simultaneously and efficiently was one of the major motivations behind the work that will be presented in Chapter 6.

The simulations from L- and I-mode which most closely matched the heat fluxes were then used to compare to the experimentally measured electron temperature fluctuations and perturbative thermal diffusivity. These simulations are referred to as the ion heat flux-matched simulations.

While particle transport is not the focus of this validation study, or of this thesis in general, it is worth mentioning that a trace impurity species ($Z=18$, $A=40$) was introduced into the ion heat flux matching simulations in order to calculate the diffusion and convection in L- and I-mode, as done in Reference [129]. The simulations show

similar levels of diffusion, convection, and peaking in both L- and I-mode, which is consistent with experimental observations that I-mode has particle confinement similar to that in L-mode [89]. Future, more detailed comparisons of experimental and simulation particle transport will have to take into account recent advances in analysis techniques for constraining experimental particle transport coefficients [50].

Comparison of Heat Fluxes to GYRO Simulations

Figure 5-4 shows the results of the electron and ion heat flux comparisons in both L- and I-mode time periods. For the experimental results, blue is L-mode and red is I-mode, as before. The GYRO results are shown in black. The shaded regions represent uncertainties, calculated by propagating the uncertainties in the inputs to TRANSP through the governing equations, as described in Chapter 2. The simulations in this figure are those for which a/L_{Ti} was scanned within experimental uncertainty in order to most closely match the ion heat flux at all radii. The input a/L_{Ti} was reduced by 1% in L-mode and by 29% in I-mode. As stated above, scanning the electron temperature gradient had little impact on the electron heat flux, and therefore did not improve agreement with experiment.

In both L- and I-mode, the GYRO simulations are able to match the experimental ion heat flux to within the experimental uncertainty over the entire radial range that was simulated. On the other hand, the GYRO simulations robustly under-predict the electron heat flux at all radii in L-mode, and over the majority of the domain in I-mode. GYRO can match the electron heat flux between ρ_{tor} of 0.8 and 0.88 in I-mode.

In order to confirm the accuracy of these global ion-scale simulations, local nonlinear GYRO simulations were performed at a few radii in the global simulation domain, and in all cases exhibit agreement within 10% of the global simulations. In particular, the results of local GYRO runs at $\rho_{tor} = 0.85$ are shown in Figure 5-4 as the grey

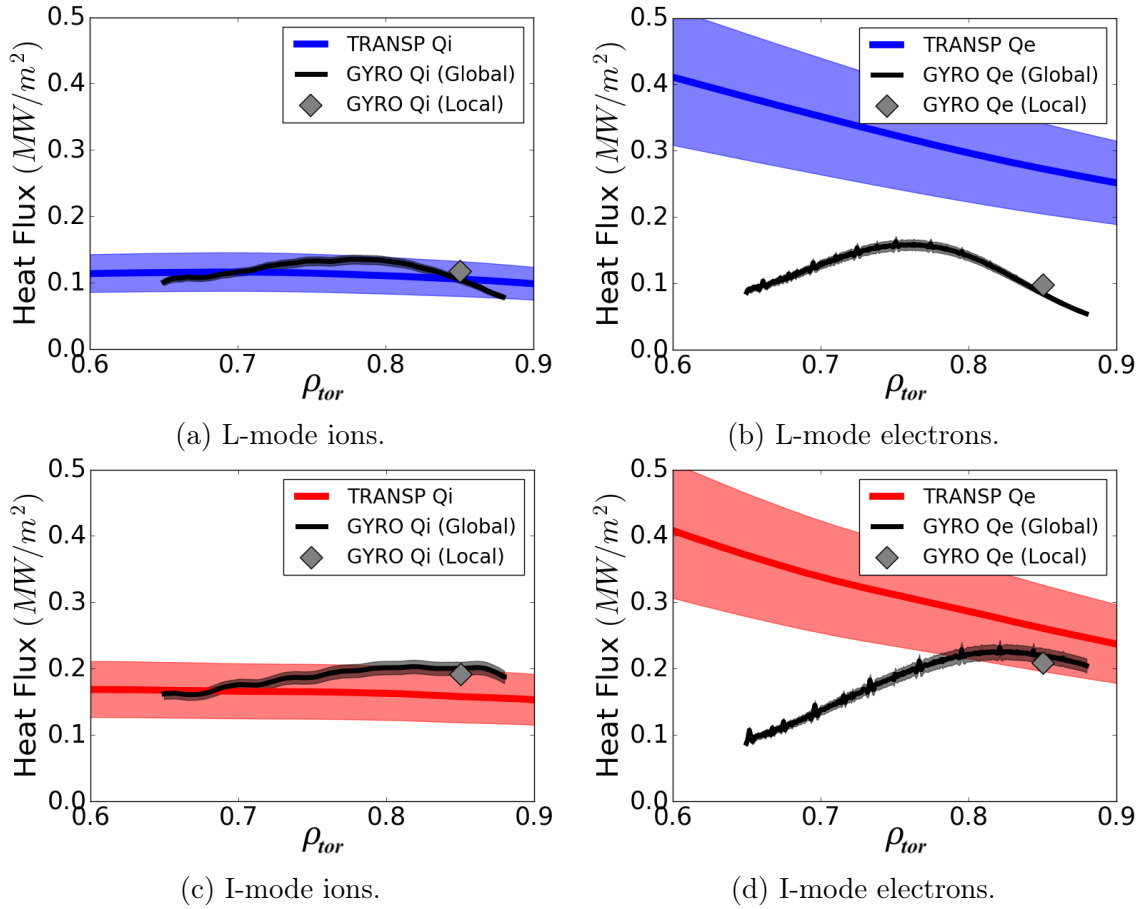


Figure 5-4: Comparison of experimental and simulated electron and ion heat fluxes in the L- and I-mode time periods. Ion heat fluxes are shown on the left and electron heat fluxes on the right. The L-mode comparison is shown on the top (with the experiment in blue) and the I-mode comparison is shown on the bottom (with the experiment in red). Experimental uncertainty is represented by the shaded regions. Global GYRO results are shown in black, with uncertainty represented as the shaded region. Gray diamonds represent local GYRO simulations, used as a check of the global results. Figure originally from Reference [100].

diamonds, and agree closely with the global simulations.

As has been discussed earlier in this thesis, high-k electron temperature gradient (ETG) turbulence can also contribute to the electron heat flux. The linear analysis shown in Figure 5-3 reveals this turbulence to be unstable. Even though past work has shown that linearly adding ion- and electron-scale turbulence is not particularly rigorous (motivating true multi-scale simulations) [14], for the sake of completeness local nonlinear electron-scale GYRO simulations were performed for both L- and I-mode at $\rho_{tor} = 0.8$ in order to investigate the effect of pure ETG turbulence. The electron-scale simulations box size was approximately 9 by 6 ρ_s in the radial and binormal directions, with a grid spacing of $2\rho_e$. These simulations captured short wavelength turbulence up to approximately $k_\theta\rho_s \approx 56$. All other inputs were the same as listed in Table 5.2.

These electron scale simulations show that pure ETG turbulence drives approximately 10 - 20% of the experimental electron heat flux, which is non-negligible, but not nearly enough to resolve the discrepancy between the ion-scale simulation results and the experimental values in either L- or I-mode. More quantitatively, in L-mode at $\rho_{tor} = 0.8$, the experimental heat flux was approximately 0.31 MW/m². The ion-scale simulation value, shown in Figure 5-4 was approximately 0.14 MW/m². The electron-scale simulation heat flux was approximately 0.03 MW/m². Even added together, the ion- and electron-scale simulation electron heat fluxes under-predict the experimental electron heat flux by more than 45% at this radius. The I-mode simulation shows similar results. Again, this process of linear addition is included only for completeness, and is not generally equivalent to running true multi-scale simulations.

The robust under-prediction of the electron heat flux observed in both L- and I-mode is consistent with a broader trend of electron heat flux under-prediction on Alcator C-Mod L-modes that is well documented [14, 126]. The only previous set of GYRO simulations of an Alcator C-Mod I-mode found a similar trend [97]. In that

case, a local simulation at $\rho_{tor} = 0.6$ under-predicted the electron heat flux, while a local simulation at $\rho_{tor} = 0.8$ agreed within experimental uncertainty.

In the L- and I-mode plasmas considered here, ion-scale GYRO simulations are already unable to match experimental parameters, even before adding additional validation constraints. To the extent that validation is the question of whether or not the experiment and simulation agree within uncertainty, in this particular case, the answer is already that they do not. For further physics insight, however, it is certainly worthwhile to compare the other validation constraints, especially since they have already been measured. This work also contributes to a better understanding of how different validation constraints interact, and how one must use as many as possible to get a complete picture of the simulation validity.

Comparison of Electron Temperature Fluctuations

In order to properly compare the electron temperature fluctuations predicted by GYRO, or any gyrokinetic simulation, to experimental measurements, one must apply a synthetic diagnostic to the outputs of the code [130, 131]. The synthetic diagnostic accounts for the physical limitations of the real diagnostic, which only imperfectly measures the true point fluctuations in the electron temperature of the plasma. Specifically, the synthetic diagnostic includes the effects of a finite beam width and emission volume (which is assumed to be a Gaussian with $L_r = 1.2$ cm and $L_z = 0.64$ cm [76]), the k-sensitivity of the physical diagnostic, and the effects of plasma rotation on the measured spectrum. One can also think of the synthetic diagnostic as mapping the diagnostic to the simulation phase space. The details of the specific synthetic diagnostic used here are described in Reference [7].

The comparisons between the measured and simulation fluctuation spectra at $\rho_{tor} = 0.8$ for L- and I-mode are shown in Figure 5-5 (plotted in terms of absolute fluctuation cross-power, instead of coherence, as this is the most appropriate com-

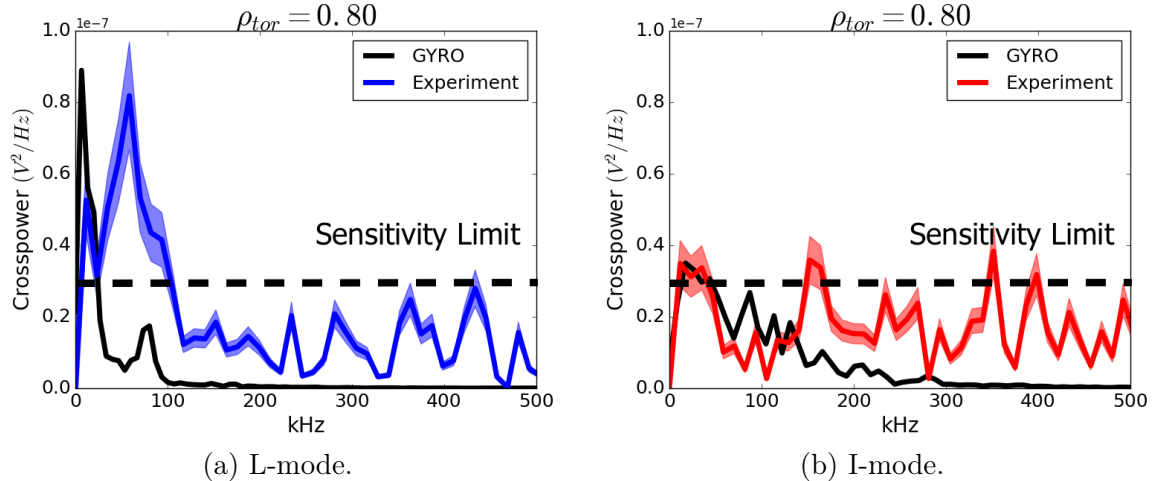


Figure 5-5: Comparison of experimentally measured and simulated electron temperature fluctuation spectra in the L-mode (a) and I-mode (b) cases. The L-mode experimental data is shown in blue and the I-mode data is shown in red. Shaded regions represent experimental uncertainty. The results of the GYRO simulations, after analysis with a synthetic diagnostic, are shown in black. The sensitivity limit is shown as a dotted black line. Originally from Reference [100].

parison with the code). As stated above, the experimental integrated temperature fluctuation level at this radius is $\tilde{T}/T = 0.81\%$ in L-mode and $\tilde{T}/T < 0.40\%$ (below the sensitivity limit) in I-mode. GYRO predicts a value below the 0.40% sensitivity limit in both L- and I-mode. The simulation thus significantly under-predicts the electron temperature fluctuation level in L-mode (by at least a factor of two). In I-mode, since both values are below the sensitivity limit, the most that can be said is that the simulation is not inconsistent with the experiment.

The L-mode result contrasts with some previous work with L-mode plasmas on Alcator C-Mod, in which GYRO agreed quite closely with the measured temperature fluctuation level even when the electron heat flux was under-predicted by the code [7]. That study in particular revealed the importance of including all of the validation constraints, not just those that are measured most directly, as fortuitous agreement between code and experiment can occur for any one of the constraints. Having many constraints, based on independent measurements of the plasma, reduces the chances of such fortuitous agreement considerably.

One may note that in the L-mode spectrum shown in Figure 5-5 there is a feature that is nearly the correct amplitude as the experimental feature, but at considerably lower frequency. Even though rotation does not influence the total fluctuation level, it can influence the shape of the spectrum, so the experimental rotation was varied within uncertainty (through the synthetic diagnostic) in order to see if doing so would improve the agreement between the experimental and simulated spectra. Such variations did not, however, result in significantly better matches.

Comparison of Perturbative Diffusivity

The final validation constraint applied in this study is the perturbative thermal diffusivity. As described in Chapter 3, the perturbative diffusivity was measured experimentally with partial sawtooth heat pulses. The simulation perturbative diffusivity, on the other hand, is calculated by performing a series of gyrokinetic simulations, in which one scans the input electron temperature gradient in order to map out the flux-gradient relationship. One begins with the heat-flux matched simulation (or the simulation that most closely matches the experimental heat fluxes). One then scans the input electron temperature gradient up and down 12%, as this approximates the temperature gradient change induced by the partial sawtooth heat pulses. If the relationship between gradient and heat flux deduced by this method is clearly non-linear, one may also wish to perform additional simulations at 6% (or 24%) above or below the flux-matched input electron temperature gradient. With the electron temperature gradient scan completed, one then plots the output electron heat flux against the input electron temperature gradient multiplied by the density. This gives a slope on this plot the proper units of m^2/s . One then fits the set of simulations, and calculates the slope of the fit at the heat flux matched simulation. This slope is the value of the perturbative thermal diffusivity for the simulation. To be explicit (in response to an occasional misunderstanding), heat pulses are not directly simulated in GYRO. The




Figure 5-6: Comparison of the experimentally measured and simulated perturbative thermal diffusivity. The electron heat fluxes, Q_e , calculated from an electron temperature gradient scan of the GYRO simulations are shown as the squares. L-mode is in blue, I-mode is in red. The circled points are the base case ion heat flux-matched simulations and the other points are scans of a/L_{Te} around these simulations. A linear fit to the scan of simulations is shown in the solid lines, the slope of which is the GYRO perturbative thermal diffusivity, χ_{GYRO}^{pert} . The dotted lines represent the experimentally measured χ_{Exp}^{pert} . One should compare the slopes of the solid and dotted lines. Figure modified from Reference [100].

method described here has been used before to compare simulation to experimental perturbative diffusivity (measured with modulated ECH) [63].

Since gyrokinetic simulations typically take the normalized electron temperature gradient scale length, a/L_{Te} , as an input, rather than the temperature gradient directly, one can write:

$$\chi_{sim}^{pert} = \frac{a}{n_e T_e} \frac{\partial Q_e}{\partial (a/L_{Te})} = \frac{1}{n_e} \frac{\partial Q_e}{\partial \nabla T_e} \quad (5.1)$$

where χ_{sim}^{pert} is the simulation (GENE, GYRO, TGLF, etc.) perturbative diffusivity, a is the plasma minor radius, n_e is the local plasma density, T_e is the local plasma temperature, and Q_e is the electron heat flux.

This process is illustrated in Figure 5-6, which shows the results of the GYRO electron temperature gradient scans, as well as the experimentally measured pertur-

bative diffusivity. In this particular case, simulations with a/L_{Te} reduced by 12%, raised by 12% and raised by 25% were performed. In this plot, as before, blue is L-mode and red is I-mode. The boxes represent the GYRO simulations and the solid lines are linear fits to these simulations. The circled simulations are the ion heat flux-matched simulations shown above. The dashed lines show the experimentally measured perturbative diffusivity (which is a slope on this plot). The simulation heat flux values are averaged over the GYRO simulation domain, in order to replicate as closely as possible the nature of the experimental measurement.

In L-mode, GYRO predicts a perturbative diffusivity of $0.4 \text{ m}^2/\text{s}$, compared to the experimentally measured $4.0 \pm 0.6 \text{ m}^2/\text{s}$. In I-mode, GYRO predicts $1.0 \text{ m}^2/\text{s}$, compared to the experimentally measured $4.3 \pm 0.9 \text{ m}^2/\text{s}$. GYRO clearly under-predicts the perturbative thermal diffusivity in both L- and I-mode, far outside of experimental uncertainty. This completes the set of validation constraints that are applied to these L- and I-mode plasmas.

Conclusions of this Validation Study of GYRO on Alcator C-Mod L- and I-mode Plasmas

To summarize the results of this validation study, global, nonlinear, ion-scale GYRO simulations are able to match the ion heat flux in both the L- and I-mode plasmas considered in this study across the simulation domain of $\rho_{tor} = 0.65 - 0.88$. On the other hand, GYRO under-predicts the electron heat flux in both L- and I-mode, with the exception of the outer region of the I-mode, between ρ_{tor} of 0.8 and 0.88. The linear addition of electron scale heat flux does not resolve this discrepancy. GYRO also under-predicts the electron temperature fluctuation level in L-mode. In I-mode, both the experimental measurement and simulated value are below the experimental uncertainty, so all that can be said is that the simulation is not inconsistent with the experiment. Finally, GYRO significantly under-predicts the perturbative thermal

diffusivity in both L- and I-modes, with an entire order of magnitude under-prediction in L-mode.

These plasmas thus represent an instance where ion-scale GYRO simulations decidedly disagree with several experimental validation constraints. As has been discussed earlier in this thesis, and will be discussed in more detail later, it is possible that running multi-scale simulations would help resolve these discrepancies. The results of ion- and multi-scale simulations will be compared directly for a different plasma discharge in the next study presented in this thesis. In addition, as has been suggested earlier in this section, it may also be that one must vary more than just the ion temperature gradient within experimental uncertainty in order to match the experimental constraints. Both of these possibilities serve as motivation for the work with VITALS and TGLF presented in Chapter 6.

It is also possible that the GYRO simulations in this study left out other important physical effects that may have an important impact on the results. For example, several relatively recent studies with turbulent transport models have shown that fast ions can have a significant impact on both the experiment itself and the outputs of gyrokinetic simulations [28, 29]. Since Alcator C-Mod uses minority ICRH as its primary heating source, there will always be some population of fast minority ions, which then collide with the bulk plasma for heating. It may be that including the effects of fast ions is significant in these simulations, though such work is beyond the scope of this thesis.

Another question that one may raise concerning the results of this validation study is whether another turbulent transport model, such as GENE, would more closely agree with experimental measurements, especially considering the fact that GENE is used in the last two validation studies in this chapter. Though not done directly as part of this thesis, GENE and GYRO were compared to one another in Reference [132] using exactly the experimental conditions presented in this validation study.

This benchmarking study (comparison of different codes to one another, as opposed to validation of a model against experiment)² found that in fact local, nonlinear, ion-scale GENE and GYRO simulations disagree when the exact same inputs are used. GENE and GYRO predict the same ratio of ion and electron heat flux, but GENE predicts that both heat fluxes are nearly three times larger than GYRO's predictions. This significantly improves the agreement with experimental electron heat flux, but leads to a large overprediction of the ion heat flux. Electron temperature fluctuations and perturbative thermal diffusivity were not compared to the outputs of those particular GENE simulations.

This benchmarking study reveals that different implementations of the gyrokinetic model may indeed give different results, and that each implementation must then be validated separately. In general, benchmarking involves extensive simulation input and parameter scans and meticulous comparisons of how inputs are expressed in different codes. Deep investigations of how physical processes are implemented computationally are often required in order to resolve discrepancies. For these reasons, benchmarking of gyrokinetic codes is beyond the scope of this thesis. Please see References [133, 134, 135, 136, 137, 138, 139, 140, 141, 142, 143] for examples of verification and benchmarking studies of GENE and GYRO, as well as various other turbulent transport models.

Now that one complete validation study has been presented, this chapter will next describe work in which the perturbative diffusivity was added as a validation constraint to several sets of simulations that had already been performed. This next study also addresses the question of whether the inclusion of multi-scale effects can resolve the discrepancy between experimentally measured and ion-scale simulation perturbative diffusivity by presenting a direct comparison.

²Note that these types of comparisons are sometimes also called verification, though technically verification is the process of determining whether or not a code has correctly implemented a model.

5.1.2 Validation of Ion- and Multi-Scale GYRO Simulations with Perturbative Diffusivity

As was described in Chapter 2, one of the most recent developments in gyrokinetic validation is the recognition that in some instances, ion-scale simulations do not sufficiently describe the plasma behavior, and that linearly adding the results from ion- and electron-scale simulations does not properly capture the interactions between the two scales. This was exactly the case in the validation study of L- and I-mode just presented. In many of these instances, multi-scale simulations, which simultaneously capture ion-scale and electron-scale dynamics, are required in order to properly capture the true plasma behavior. See References [10, 11, 12, 13, 14, 15, 16, 17, 55, 126, 144] for a number of recent studies on this topic. In light of this, there is considerable motivation to find validation constraints which can easily distinguish the performance of ion- and multi-scale simulations, and can help further determine when multi-scale simulations are necessary. This work will show that the perturbative diffusivity is one such measurement.

The validation study presented here, originally published in Reference [55], focuses solely on the use of the perturbative diffusivity as a validation constraint, and is based off of simulations that were performed previously as part of other studies. The original gyrokinetic simulations were performed as part of References [12, 97] but the addition of the perturbative diffusivity as a constraint and the resulting conclusions were a contribution of this thesis work. As such, this section will primarily describe how the perturbative diffusivity was used as a constraint in these validation efforts.

The GYRO simulations in this validation study were local, non-linear, flux-tube simulations of Alcator C-Mod discharges 1101209029 [97] and 1120221012 [12]. The first discharge (1101209029) has only ion-scale simulations available, but contains both L- and I-mode portions of the discharge. The second discharge (1120221012)

Table 5.3: Table from Reference [55]. Comparison of the experimental and GYRO perturbative diffusivities from discharges 1101209029 and 1120221012. Discharge 1101209029 has both L- and I-mode time periods. Discharge 1120221012 has only one experimental case, but has both ion- and multi-scale simulations. All GYRO results are from [12] and [97].

Shot	Confinement	GYRO Model	$\chi_{Partial}^{pert}$ (m^2/s)	χ_{GYRO}^{pert} (m^2/s)
1101209029	L-Mode	Ion-scale	1.7 ± 0.3	0.3
1101209029	I-Mode	Ion-scale	2.0 ± 0.4	0.8
1120221012	L-Mode	Ion-scale	1.6 ± 0.4	0.2
1120221012	L-Mode	Multi-scale	1.6 ± 0.4	1.4

is L-mode, and both ion- and multi-scale simulations have been performed. In both cases, the inputs to GYRO were taken directly from experiment. Please see References [12, 97] for details of the specific gyrokinetic simulations.

In both cases, the input ion temperature gradient (expressed as the normalized ion temperature gradient scale length a/L_{Ti}) was varied within the experimental uncertainty in order to match the experimentally inferred heat flux (from TRANSP, as described in Chapter 2). The simulations presented here are the ion heat flux matching simulations from each respective discharge. References [12, 97] discuss comparisons with heat fluxes in more detail.

As was described in the last validation study, the simulation perturbative diffusivity is calculated by performing a series of GYRO simulations in which one scans the input electron temperature gradient up and down by 12%, which approximates the temperature gradient change induced by the partial sawtooth heat pulses. One then fits the flux-gradient curve of these simulations, and takes the slope of the fit at the heat flux matched simulation in order to calculate the simulation perturbative diffusivity.

This process was performed for the four simulation cases described above: ion-scale L-mode in 1101209029, ion-scale I-mode in 1101209029, ion-scale L-mode in

1120221012, and multi-scale L-mode in 1120221012. The results of all of these simulation scans, as well as the experimentally measured values (measured with partial sawtooth heat pulses as described in Chapter 3), are shown in Table 5.3. In all three ion-scale simulation cases, the GYRO simulations significantly under-predict the experimentally measured perturbative diffusivity, sometimes by nearly an order of magnitude. This is similar to the result from the validation study presented earlier in this chapter. In stark contrast, the multi-scale simulation of discharge 1120221012 agrees with the experimental perturbative diffusivity to within experimental uncertainty.

While Table 5.3 only contains one direct comparison of ion- and multi-scale simulations, the fact that the ion-scale simulation robustly under-predicts the experimentally measured perturbative diffusivity, while the multi-scale simulation agrees within uncertainty, is strong evidence that in at least some instances, the inclusion of multi-scale effects can resolve the discrepancy between experimentally measured validation constraints and the outputs of ion-scale simulations. This is consistent with the hypothesis presented in the last validation study, that multi-scale effects may at least in part be responsible for the discrepancy between experiment and simulation. This result also motivates the work presented in Chapter 6, which seeks to determine when multi-scale effects are important in turbulent transport simulations, and when ion-scale models are sufficient.

In addition, it is interesting that even though the ion-scale simulations of both the L- and I-mode time periods in 1101209029 were able to match the experimental heat fluxes (at $\rho_{tor} = 0.8$) [97], they were unable to match the experimental perturbative diffusivity. This would seem to indicate that even though it is possible to adjust the simulation input gradients to achieve a heat flux match, the underlying response of the plasma turbulence to the input gradients is not properly captured by the simulations. This is a perfect example of what is known as ‘fortuitous agreement,’ where the

simulations agree with the experiment, but are in reality not capturing all of the relevant physics. It is precisely to avoid fortuitous agreement that one must use as many validation constraints as are available. In this sense, the perturbative diffusivity seems to be a powerful constraint in determining whether or not the simulations are properly capturing the plasma behavior.

While the study presented here and the first one in this chapter combined represent only five experimental cases (four with only ion-scale simulations and one with both ion- and multi-scale simulations), the results may lead one to believe that multi-scale simulations are always necessary to properly capture the plasma perturbative diffusivity. This turns out not to be the case, as will be demonstrated later in this chapter in Section 5.2. The question of when multi-scale simulations are necessary, and when ion-scale simulations sufficiently capture all of the relevant plasma behavior, will continue to arise through the rest of this thesis, and is the central motivation for the work presented in Chapter 6

Conclusions from GYRO Validation Studies on Alcator C-Mod

This section has presented two validation studies with GYRO on Alcator C-Mod. One was a complete validation study, focusing on the differences between L- and I-mode plasmas, while the other added perturbative diffusivity as a validation constraint to several previously existing sets of simulations. Taken together these results show that in many cases ion-scale simulations do not adequately model the plasma behavior, disagreeing with validation constraints outside of experimental uncertainty. These studies highlight the importance of including many validation constraints, in order to avoid fortuitous agreement. Finally, these results suggest that multi-scale effects may resolve some of the discrepancies observed between ion-scale simulations and experimental measurements. The next section in this chapter continues this work, using the same set of validation constraints on the ASDEX Upgrade tokamak. Validation

efforts on a second machine expand the accessible parameter space, and allow one to eventually build a more complete picture of the physics behind the agreement or disagreement of ion-scale simulations with experimental constraints.

5.2 Validation of GENE on ASDEX Upgrade

This section will present the results of two validation studies with the gyrokinetic code GENE on ASDEX Upgrade. In both cases, the GENE simulations are local, nonlinear, and ion-scale. As described in Chapter 2, ASDEX Upgrade is a fairly different machine to Alcator C-Mod, thus allowing access to different plasma parameter ranges and different turbulent regimes. As such, validation on ASDEX Upgrade complements validation on Alcator C-Mod.

GENE was chosen, instead of GYRO, primarily for practical reasons. GENE has been run repeatedly on ASDEX Upgrade in past work, while GYRO has not been run on ASDEX Upgrade. For this reason, there is no expertise in running GYRO on ASDEX Upgrade, and the goal of this thesis was not to work on the logistics of running simulations, but on the validation of turbulent transport models. In addition, benchmarking of different gyrokinetic codes against one another is beyond the scope of this thesis, though some contributions were made to a benchmarking study of GENE and GYRO performed on the L- and I-mode discharges described in the previous section. This benchmarking work is presented in Reference [132]. In addition, Chapter 6 will work to resolve the issue of running different codes on different machines by running TGLF on both Alcator C-Mod and ASDEX Upgrade discharges.

Both validation studies in this section are based on the same experimental discharge on ASDEX Upgrade, 33585, though at different radii and using different validation constraints. Note also that this same discharge is one of those considered in the multi-machine, multi-discharge study in Chapter 6. The first validation study, presented in Subsection 5.2.2, was performed entirely as part of this thesis, and was

originally presented in Reference [64]. It focuses on an inner radius of $\rho_{tor} = 0.49$ and employs heat fluxes and perturbative thermal diffusivity as validation constraints.

The second validation study, presented in Subsection 5.2.3 was not performed directly as part of this thesis, but this thesis work contributed the power balance analysis and the CECE hardware and analysis techniques described in Chapter 4. This study focused on an outer radius of $\rho_{tor} = 0.75$, and included heat fluxes, electron temperature fluctuation spectra and radial correlation lengths, and density-temperature fluctuation phase (n-T phase) measurements as validation constraints (which are briefly explained below). Results from this study were originally presented in Reference [107].

This section first briefly describes the experimental discharge, before then describing both validation studies.

5.2.1 Experimental Conditions of ASDEX Upgrade Discharge 33585

ASDEX Upgrade discharge 33585 was an L-mode plasma with a toroidal field on axis of $B_t = 2.5$ T, plasma current $I_p = 1.0$ MA, and central density $n_{e,0} = 2.6 \times 10^{19}$ m⁻³. It was heated with 0.7 MW of electron cyclotron heating (ECH), which was deposited at $\rho_{tor} \approx 0.2$. Neutral beam blips were also included in order to allow ion temperature and rotation measurements with charge exchange recombination spectroscopy. The plasma was stationary for 2 seconds of flattop operation, from 2 to 4 seconds, which is the period for which all of the following analysis was carried out. This plasma contained partial sawteeth, allowing for perturbative diffusivity measurements.

All of the experimental measurements of this discharge were made with the diagnostics described in Chapter 2. Recall that the uncertainties in the plasma profiles measured in ASDEX Upgrade are approximately 15% in T_e , 20% in T_i , 10% in n_e ,




Figure 5-7: Temperature and density profiles from ASDEX Upgrade discharge 33585, averaged from 2 to 4 seconds. For the temperature, red is for electrons and blue is for ions. Also shown are the normalized gradient scale lengths. Figure modified from Reference [107].

and 50% in v_{rot} . These fitted profiles are then used to calculate the normalized gradient scale lengths, a/L_y , which are defined as above, except that in GENE, a is defined as the square root of the edge toroidal flux divided by π times the reference toroidal magnetic field, $a = \sqrt{\Phi_{edge}/\pi B_{ref}}$ [102]. This is generally equal to the plasma minor radius, which was the definition of a used above. The resulting uncertainties in the normalized gradient scale lengths are: 20% in a/L_{Te} , 30% in a/L_{Ti} , and 30% in a/L_{ne} . The measured plasma profiles and normalized gradient scale lengths for ASDEX Upgrade discharge 33585 are plotted in Figure 5-7.

Note that for the second of the two validation studies in this section, the electron temperature fluctuation measurements were made on a repeat discharge, since during discharge 33585 only the first version of the ASDEX Upgrade CECE system described in Chapter 4 was available. Discharge 34626 was an exact repeat of discharge

33585, though with the new CECE hardware (the second system described in Chapter 4) and therefore significantly expanded measurement capabilities. Specifically, though temperature fluctuation measurements were available in the first discharge, radial correlation length and n-T phase measurements were not. These measurements were enabled by the fine comb of channels on the upgraded CECE system, which directly allows for radial correlation length measurements, and makes the alignment of radiometer and reflectometer channels required for n-T phase measurements considerably easier. Note that discharge 34626 is one of the discharges for which the radial temperature fluctuation profile is shown in Figure 4-12.

5.2.2 Inner Radius Validation Study Perturbative Thermal Diffusivity Measurements

Consider first the inner radius study, at $\rho_{tor} = 0.49$, which was performed entirely as part of this thesis. This validation study was originally published in Reference [64], and used the electron heat flux, ion heat flux, and electron perturbative thermal diffusivity as validation constraints. This section will first describe the inputs to the simulations used in both this and the next validation studies, before then going into both the linear stability analysis results and nonlinear results. The results of linear stability analysis will focus primarily on the interaction of ion- and electron-modes at long wavelength in this plasma. The nonlinear results will compare heat fluxes and the perturbative diffusivity, as well as introducing a validation metric, which summarizes the agreement between experiment and simulation with a single number. The validation study presented in this section represents the first example of an ion-scale gyrokinetic simulation that agrees with the experimentally measured perturbative diffusivity.

Some commentary on the results of just this study will be given at the end of this subsection, but a more thorough discussion involving all four validation studies from

this chapter will be given at the end of the chapter.

GENE Simulations

The GENE simulations in these validation studies were local, ion-scale ($k_y \rho_s < 2.3$), where k_y is the binormal wavenumber and ρ_s is the ion Larmor radius evaluated at the electron temperature, and nonlinear. These simulations included electromagnetic effects, used a realistic electron to ion mass ratio, included the effects of impurities through an effective charge Z_{eff} (to be discussed in further detail later), and treated collisions with a linearized Landau-Boltzmann operator. The simulations utilized the real experimental magnetic equilibrium, as reconstructed with the CLISTE code [48].

These nonlinear simulations included two gyrokinetic species (electrons and main deuterium ions), 128 grid points in the radial direction, 48 bi-normal modes (corresponding to 96 grid points due to the hermiticity of the Fourier coefficients for real quantities [102]), 24 grid points in the parallel direction, 48 grid points in parallel velocity, and 16 grid points in magnetic moment (for a resolution of $128 \times 96 \times 24 \times 48 \times 16$ grid points). The binormal wavenumber $k_y \rho_s$, which determines whether the simulations are ion-, electron-, or multi-scale, ranged from 0.048 to 2.3, corresponding to a minimum finite toroidal mode number of 5. Keeping modes up to $k_y \rho_s = 2.3$ ensures that no ion-scale dynamics are lost, even if some modes extend to slightly shorter wavelength than the ion Larmor radius. The simulation box extended $113.5 \rho_s$ in the radial direction, up to $3v_{th,j}$ in the parallel velocity (where $v_{th,j} = \sqrt{2T_{0j}/m_j}$), and up to $9T_{0j}/B_{ref}$ in the magnetic moment. In both of the validation studies described here, convergence checks were performed on the simulation resolution and box size (except in bi-normal modes, as this would begin to extend into the electron scale) in order to ensure that the results were numerically robust.

Table 5.4 summarizes the experimental inputs used in the GENE simulations for both validation studies described in this section. Note that some parameters that

Table 5.4: Inputs to the GENE simulations at both the inner (Subsection 5.2.2) and outer (Subsection 5.2.3) radii. The magnetic shear is \hat{s} and the $E \times B$ shearing rate is $\gamma_{E \times B}$. All values are the nominal experimental values, except for the values of the normalized temperature gradient scale lengths in parentheses, which are for the heat flux matched cases of each validation study (a/L_{T_i} down by 19% and a/L_{T_e} down by 12% for the inner radius study, and a/L_{T_i} unchanged and a/L_{T_e} down by 13% for the outer radius study).

Quantity	Inner Radius Study	Outer Radius Study
ρ_{tor}	0.49	0.75
q	1.42	2.28
\hat{s}	0.728	1.923
$\gamma_{E \times B} [a/c_s]$	0.0042	0.0168
a/L_{T_i}	1.59 (1.29)	5.11
a/L_{T_e}	3.27 (2.88)	2.65 (2.31)
a/L_n	0.926	1.46
T_e [keV]	1.57	0.60
$n_e [10^{19} \text{ m}^{-3}]$	2.14	1.53
T_i/T_e	0.41	0.65
Z_{eff}	1.6	1.6
β_e [%]	1.96	5.39
ν_c	0.000209	0.000952
$\nu_{ei} [a/c_s]$	0.081	0.370
R_{axis} [m]	1.65	1.65
r_{minor} [m]	0.325	0.47
$B_{ref} = B_0$ [T]	2.626	2.626
$L_{ref} = a$ [m]	0.652	0.652
c_s [km/s]	274.16	170.30

don't depend on the radial location of the plasma are the same in both studies. All of these values are the nominal experimental values, except for those values in parenthesis, which represent variations of the input temperature gradients within experimental uncertainty in order to best match the validation constraints.

Linear Stability Results

Consider first the results of linear stability analysis performed on the discharge in question. Figure 5-8 (a) shows the linear stability analysis for the nominal experimental gradients, and (b) shows the linear stability analysis for a case where the temperature gradients have been adjusted within experimental uncertainty to better match the nonlinear heat flux validation constraints (a/L_{Ti} reduced by 19% and a/L_{Te} reduced by 12%). The nonlinear heat flux results will be described in greater detail below, but the linear results from the 'flux-matched' simulation are shown here for comparison. In this figure, blue points represent modes with a positive real frequency, indicative of modes moving in the ion drift direction. Red points represent modes with a negative real frequency, indicative of modes moving in the electron drift direction. Note that the frequency sign convention is opposite for the GENE and GYRO codes. Finally, diamonds show the dominant mode at a given wavenumber, and circles show the largest subdominant mode.

While one can analyze various aspects of the linear stability results, perhaps the most pertinent observation in Figure 5-8 is that in both cases the plasma is mixed mode at low wavenumber, particularly in the nominal case, where ion- and electron-modes alternate as the dominantly growing mode. In the nominal case, the TEM-like mode is dominant throughout the low-k region, except for around $k_y \rho_s \approx 0.9$ where the ITG-like mode briefly becomes dominant. In the flux-matched case (with a/L_{Ti} reduced by 19% and input a/L_{Te} reduced by 12%) the electron (TEM) mode is always dominant, though the ion (ITG) mode growth rate nearly equals that of the electron

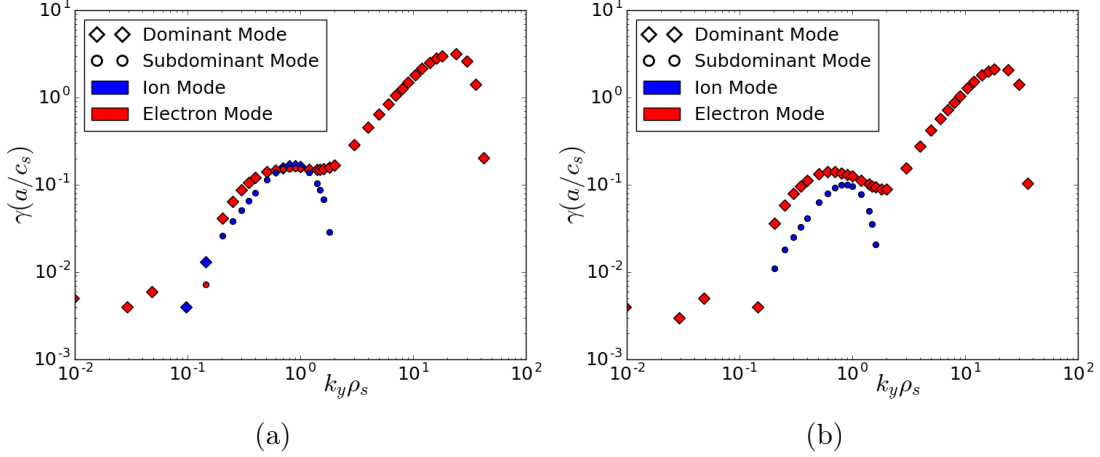


Figure 5-8: Linear stability results at $\rho_{tor} = 0.49$ of ASDEX Upgrade discharge 33585. Blue points have positive real frequency (the ion diamagnetic direction) and red points have negative real frequency (the electron diamagnetic direction). (a) shows the linear growth rates of the dominant and secondary modes for the nominal experimental parameters and (b) shows the results for the heat flux matched case (a/L_{Ti} reduced by 19% and a/L_{Te} reduced by 12%). Originally from Reference [64].

mode at around $k_y \rho_s \approx 1.0$. This strongly mixed mode character of the plasma foreshadows the unintuitive nonlinear dependence of the electron and ion heat fluxes on both temperature gradients that are discussed below.

The results of linear stability analysis were also used to justify running simulations with a Z_{eff} parameter rather than with a third gyrokinetic impurity species. In general, a simulation with a Z_{eff} parameter captures only some of the effects of impurities in a plasma, such as the effect on collisionality. It does not, however capture the effect of dilution (the discrepancy between electron and main ion density) or instabilities caused by the impurity ions themselves. In some cases, the Z_{eff} parameter does sufficiently capture the impact of impurities in the plasma, especially in cases where there is not a particularly large impurity content, and where the plasma dilution is low. On the other hand, one must sometimes run a simulation with a third gyrokinetic species in order to properly model the plasma turbulence (as was done in the Alcator C-Mod GYRO simulations in Section 5.1). The downside of doing so is that including a third gyrokinetic species considerably increases the computational

resources required by the simulation.

Linear stability results for the nominal experimental inputs were compared between simulation cases run with Z_{eff} and with a third gyrokinetic species in order to determine the feasibility of saving computational time by using Z_{eff} in the nonlinear simulations. Since ASDEX Upgrade's plasma facing components are dominantly either made of tungsten or tungsten-coated [40], tungsten was chosen as the representative impurity. The tungsten impurity was partially ionized to +46, as is suggested in Reference [145].

The linear growth rates from the simulation with a tungsten impurity and the simulation with Z_{eff} vary by less than 6% up to $k_y \rho_s \lesssim 1.0$ (and on average by less than 5%), which is past the peak in the linear growth rates. The linear results do differ slightly more at higher wavenumber, and the average difference for all wavenumber included in the nonlinear simulations (up to $k_y \rho_s = 2.3$) is approximately 12 %. The dilution of the plasma is less than 2% (the main ion density is approximately 98.7% of the electron density).

Since the linear stability results differ only slightly up to the dominantly unstable wavenumber, it was determined that for this particular instance, using a simulation with Z_{eff} would sufficiently capture the effects of impurities in the plasma, and that it was not necessary to perform nonlinear runs with a gyrokinetic impurity species.

Nonlinear Validation Results

This study will now compare the experimentally measured heat fluxes with the results of nonlinear GENE simulations. As is standard practice in validation studies, the input temperature gradients were varied within experimental uncertainty (20% for a/L_{Te} and 30% for a/L_{Ti}) in order to try to obtain a 'heat flux matched' case, in which the experimental and simulated ion and electron heat fluxes agree within uncertainty. The existence of such a case, for which inputs are varied within uncertainty and

outputs agree within uncertainty, is an important validation result in its own right. In addition, the input scans performed during the search for such a case offer insight into the way in which the heat fluxes depend nonlinearly on the input gradients.

First, consider the manner in which one calculates the experimental and simulated heat fluxes for comparison. As described in Chapter 2, the experimental heat fluxes are calculated with the power balance transport solver TRANSP [146], giving $P_e = 0.67$ MW and $P_i = 0.21$ MW for this radius of this discharge. Recall that in this thesis, P_j refers to a total power in MW while Q_j refers to a heat flux in MW/m² ($P_j = S \cdot Q_j$, where S is the flux surface area). As described before, uncertainties in the inputs to TRANSP are propagated through the code's governing equations to obtain uncertainties of 18% for P_e and 26% for P_i .

The GENE heat fluxes are then calculated by averaging the heat flux output of the simulation (adding both electrostatic and electromagnetic contributions to the total heat flux) over several hundred a/c_s times, where c_s is the sound speed. The ratio of a/c_s gives the characteristic time of the simulation output. The simulation outputs the heat fluxes many times during each a/c_s time, so these averages typically contain several thousand data points. One must also ensure that the time averaging only begins after initial transients in the simulation have died off, and the effects of $E \times B$ shear (if it is initialized after the simulation begins) have come to equilibrium.

This averaging process is illustrated in Figure 5-9, which shows data from the heat flux matched GENE run (a/L_{Ti} reduced by 19% and a/L_{Te} reduced by 12%). In this figure, the turquoise region highlights the time period over which the simulation results were averaged. In this simulation, there is an initial transient as the heat fluxes grow, then the effects of $E \times B$ shear were activated at $t = 250a/c_s$ (the vertical dashed line). After the transient caused by the $E \times B$ shear has died out, the averaging begins at $t = 375a/c_s$ and continuing for an additional 540 a/c_s . The average heat fluxes are represented as the horizontal dashed lines. The uncertainty




Figure 5-9: Heat flux output during GENE simulation of ASDEX Upgrade discharge 33585 at $\rho_{tor} = 0.49$. This is the heat flux matched case, with a/L_{Ti} reduced by 19% and input a/L_{Te} reduced by 12%. The data is plotted as Gyro-Bohm heat flux (electrostatic plus electromagnetic) against simulation time in units of a/c_s . Electrons are in red and ions are in blue. Averaging was performed over the turquoise region, with the dashed lines representing the average heat flux. Originally from Reference [64].

on this average is calculated using the method described in Reference [147].

For the particular case shown in Figure 5-9, the average heat fluxes are $Q_e = (2.13 \pm 0.06)Q_{GB}$ and $Q_i = (0.48 \pm 0.04)Q_{GB}$, where Q_{GB} is the Gyro-Bohm heat flux [148], a theoretical prediction of how turbulent heat fluxes should scale with various plasma parameters. These values are then converted to powers (units of MW) by multiplying by Q_{GB} and the flux-surface area to get $P_e^{GENE} = 0.73$ MW and $P_i^{GENE} = 0.16$ MW. This agrees with the experimental values ($P_e^{Exp} = 0.67 \pm 0.12$ MW and $P_i^{Exp} = 0.21 \pm 0.06$ MW) within uncertainty. This averaging procedure was used on all of the simulations presented in this study, including those used to check for numerical convergence in box size and resolution. A similar procedure was used to obtain the heat fluxes in the GYRO validation studies presented earlier in this chapter, though the fact that the simulations were global somewhat complicates the process.

One interesting observation in Figure 5-9 is that even though there is a strong

transient after the initialization of $E \times B$ shear, the average heat fluxes quickly return to values that are very close to what they were before the effect was activated. This same phenomenon is observed in all of the simulations in this study, and can be explained by the relatively low value of the $E \times B$ shear in Table 5.4, which is much smaller than the peak linear growth rates shown in Figure 5-8. Experimentally, the small $E \times B$ shear is a result of a relatively flat rotation profile at the radius under consideration. One should note, however, that rotation profile measurements can be fairly uncertain, and that in cases where the $E \times B$ shear is more important, this can be a potential source of uncertainty in the final simulation outputs. There is also some disagreement within the simulation community concerning the correct way to implement the effects of $E \times B$ shear in simulations.

As was described above, it is possible to separately perform ion-scale and electron-scale simulations, calculate heat fluxes, and linearly add these results together, though past work has shown that this procedure often gives poor agreement with real multi-scale simulations [10, 14]. This was, for example, performed for the L- and I-mode validation study with GYRO on Alcator C-Mod described above. The simulations in the validation study now under consideration, however, have a very low ratio of peak growth rate at the electron scales to the peak growth rate at the ion-scales (see Figure 5-8), and so it is unlikely that electron scales will contribute significantly to the heat flux. The importance of the ratio of the electron- to ion-scale peak growth rates will be discussed in further detail in Chapter 6. For this reason, only ion-scale simulations are performed as part of this study.

In practice, the heat flux matched simulation shown in Figure 5-9 was obtained iteratively, first using the nominal experimental inputs in the simulation, and then adjusting a/L_{Ti} and a/L_{Te} within the experimental uncertainty in order to obtain better matches with the experimental heat fluxes. Table 5.5 shows the results of some of these simulations, not including some simulations that were clearly not going

Table 5.5: Summary of the nonlinear GENE simulations performed as part of this validation study. The table contains the input normalized gradient scale lengths (relative to the nominal experimental parameters), the simulation time over which averaging was performed, and the output heat fluxes. Rows are in order of decreasing ion temperature gradient and then decreasing electron temperature gradient. In Gyro-Bohm units, the experimental heat fluxes were $Q_e^{exp} = 1.97 \pm 0.35 Q_{GB}$ and $Q_i^{exp} = 0.62 \pm 0.16 Q_{GB}$, where $Q_{GB} = 0.016 \text{ MW/m}^2$. The heat flux matched simulation is highlighted in green. Originally from Reference [64].

a/L_{T_i}	a/L_{T_e}	Averaged a/c_s	Q_i/Q_{GB}	Q_e/Q_{GB}
Nominal	Nominal	110	2.50 ± 0.15	5.03 ± 0.22
-10%	Nominal	179	1.08 ± 0.08	4.41 ± 0.29
-19%	-12%	540	0.48 ± 0.04	2.13 ± 0.06
-20%	Nominal	180	0.93 ± 0.06	7.12 ± 0.55
-20%	-13%	370	0.41 ± 0.05	2.01 ± 0.24
-20%	-15%	595	0.56 ± 0.02	1.42 ± 0.08
-20%	-20%	300	1.02 ± 0.03	1.47 ± 0.03
-21%	-15%	375	0.43 ± 0.03	1.53 ± 0.07
-22%	-15%	290	0.34 ± 0.01	2.08 ± 0.05
-24%	-15%	195	0.43 ± 0.02	3.03 ± 0.11
-25%	-15%	300	0.37 ± 0.03	2.63 ± 0.24

to match the experimental heat fluxes and so were terminated before a reasonable averaging time had been completed (in order to save computational resources). Since the relationship between the heat fluxes and gradients is highly nonlinear, the iteration process can be rather non-intuitive and often requires many simulations. These scans are still, however, interesting in their own right, as they reveal more about the nature of the turbulent modes that are active in this discharge.

The first significant result is that one is indeed able to find a heat flux matched solution, as was described above. This flux matched simulation will also be the one that is compared to the perturbative diffusivity in the next section. Also interesting is the result that one must change *both* a/L_{T_i} and a/L_{T_e} simultaneously in order to match the electron and ion heat fluxes. Not only this, but one must change both

gradients to match *just* the ion heat flux, which was perhaps foreshadowed by the mixed mode nature of the plasma revealed by the linear stability analysis.

More specifically, Table 5.5 shows that the nominal experimental gradients overpredict Q_i by almost a factor of four and Q_e by more than a factor of two. The next simulation, which reduced a/L_{Ti} by 10%, reduced the overprediction to roughly a factor of two in both Q_i and Q_e . Further reducing a/L_{Ti} to 20% below the nominal value has little effect on the ion heat flux, but worsens the overprediction of the electron heat flux. This initial scan of a/L_{Ti} indicated that one would have to change *both* a/L_{Te} and a/L_{Ti} in order to match either heat flux.




Figure 5-10: GENE electron and ion heat fluxes as a function of input electron temperature gradient (all at the same ion temperature gradient), compared to the experimental value. Ion heat fluxes are in blue and electrons in red. GENE values are given by the triangles, with uncertainty represented as the error bars. The experimental values are given as the solid lines, with uncertainty represented as the shaded regions. Originally from Reference [64].

Consider changing a/L_{Te} from nominal to 20% below nominal at fixed a/L_{Ti} (20% below nominal), which is shown in Figure 5-10. In this figure, the electron channel is shown in red and the ion channel is shown in blue, with uncertainty represented as the shaded regions. Initially, reducing a/L_{Te} from the nominal value reduces P_e from

a large overprediction to roughly the experimental level (at 13% below the nominal value), agreeing within uncertainty. Further reduction brings the predicted P_e below the experimental level, but then leads to a slight increase beyond 15% below the nominal value. Interestingly, P_i shows fairly strong non-monotonic behavior as one decreases a/L_{Te} , since P_i initially decreases with decreasing a/L_{Te} and then later increases. Note that if one considers both experimental and simulation uncertainty, the simulation at -20% a/L_{Ti} and -13 % a/L_{Te} agrees with experiment, though not as well as the flux matched simulation described above.

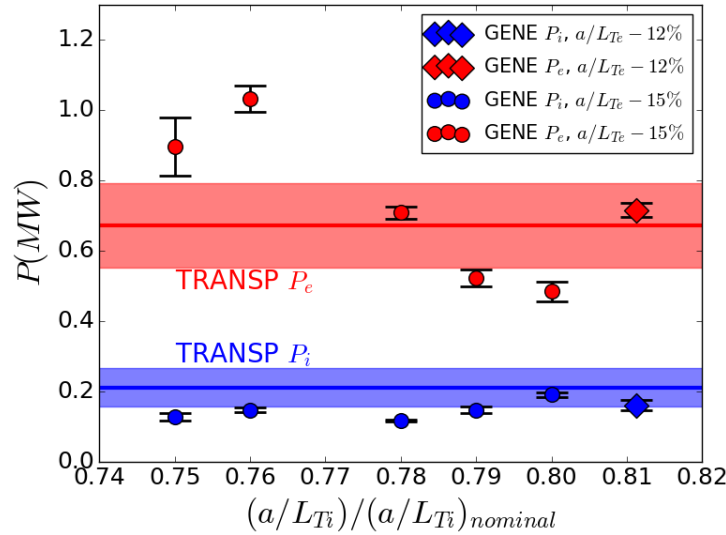


Figure 5-11: GENE electron and ion heat fluxes as a function of input ion temperature gradient (all at the same electron temperature gradient), compared to the experimental value. Ion heat fluxes are in blue and electrons in red. GENE values with $a/L_{Te} - 15\%$ are given by the triangles, with uncertainty represented as the error bars. The experimental values are given as the solid lines, with uncertainty represented as the shaded regions. The flux matched case ($a/L_{Te} - 12\%$ and $a/L_{Te} - 19\%$) is shown as the diamonds. Originally from Reference [64].

In addition, a scan of a/L_{Ti} at fixed a/L_{Te} (15% below the nominal value) was performed, and is shown in Figure 5-11 as the circles. This figure also shows the flux matched simulation as diamonds for the sake of comparison. Note that Figures 5-10 and 5-11 have different x- and y-axes. This figure reveals weaker non-monotonic behavior than for the scan of a/L_{Te} , but it does occur outside of the simulation

uncertainty. Note also that the point at $-20\% a/L_{Ti}$ and $-15\% a/L_{Te}$ is the same on both figures. This observed non-monotonic behavior of the heat fluxes with regards to the temperature gradients is consistent with the mixed-mode nature of the plasma revealed by the linear stability results.

It is worth recognizing that a few points on both of these plots agree with P_i within uncertainty, but under-predict P_e by some amount. As discussed above, it is possible that purely electron scale turbulence contributes some small portion of the electron heat flux, though the small linear growth rates at these scales suggest that this contribution would be minimal. It is therefore possible that if one linearly added electron- and ion-scale contributions some of these simulations would agree with the experiment within uncertainty. As stated before, however, linearly adding different scales is not particularly rigorous and for this reason the flux matched simulation from above is taken as the closest match to experiment.

Now that a heat flux matched simulation has been identified, the experimentally measured electron perturbative thermal diffusivity will be applied as an additional validation constraint. Using the method described in Chapter 3, the experimental perturbative diffusivity in this discharge was $\chi_{Exp}^{pert} = 7.7 \pm 2.3 \text{ m}^2/\text{s}$, averaged over the radial range of $\rho_{tor} = 0.44 - 0.54$. The radial extent of the perturbative diffusivity measurement is the reason that the GENE simulation location was chosen to be $\rho_{tor} = 0.49$.

The simulation perturbative diffusivity is calculated in the same manner as for the GYRO validation study presented earlier in this chapter, scanning the input electron temperature gradient up and down by 6% and 12% (a smaller range was used due to the smaller amplitude heat pulses on ASDEX Upgrade) and then taking the slope of the flux-gradient relationship.

In terms of time scales, the heat pulse propagates much more slowly than any turbulence time scale. Quantitatively, the sawtooth period is approximately 25 ms ,

while the a/c_s time scale of the turbulence is approximately $2.4 \mu\text{s}$. Even if the turbulence saturation time is $100 a/c_s$, there are still 100 saturation times in a single sawtooth period.




Figure 5-12: The experimental and simulation perturbative thermal diffusivities from ASDEX Upgrade discharge 33585. Red squares represent GENE simulations, showing the electron heat flux for a given a/L_{Te} input. The circled simulation is the flux matched case. The dashed red line is an exponential fit to the five GENE simulations, and the solid line is the tangent to this fit at the flux matched simulation. The turquoise line is the experimental perturbative diffusivity, with uncertainty represented as the shaded region. One should compare the solid red line with the dash-dotted turquoise line. Figure originally from Reference [64].

The result of the perturbative thermal diffusivity analysis in GENE is shown in Figure 5-12. The flux matched simulation just described is the circled point in the plot, and the experimental value of the perturbative diffusivity is the slope of the dash-dotted turquoise line going through the heat flux matched simulation (the shaded region represents the experimental uncertainty).

GENE simulations are represented as red squares on this plot and are fit with an exponential curve (the red dotted line). The slope of this fit at the flux-matched simulation (the line tangent to to the fit curve) is shown as the solid red line, which gives a perturbative diffusivity value of $8.8 \text{ m}^2/\text{s}$. The simulation value of the pertur-

bative diffusivity therefore agrees with the experimentally measured value of 7.7 ± 2.3 m²/s within uncertainty. Note that if one uses a linear fit to the GENE simulations, instead of an exponential fit, one calculates a perturbative diffusivity of 11.3 m²/s, which is slightly outside of the experimental uncertainty. This result highlights the importance of careful fitting of the simulation data when calculating the simulation perturbative diffusivity.

This agreement contrasts starkly with the results of ion-scale gyrokinetic simulations on Alcator C-Mod, such as those presented earlier in this chapter and in References [10, 14, 17, 55, 100]. In all of these past cases, ion-scale gyrokinetic simulations severely under-predicted the perturbative diffusivity, by up to an order of magnitude as in the L- and I-mode study presented earlier. This under-prediction even occurred in situations where the gyrokinetic code was able to match the experimental electron and ion heat fluxes [17], and only by running a multi-scale simulation was it possible to resolve the discrepancy.

The agreement between experiment and simulation found in this case strongly supports the theory that multi-scale effects and cross-scale coupling are very important in turbulent transport in some plasma conditions, and less important in other cases for which ion-scale simulations are sufficient. This is discussed further later in this chapter, and the question of what differentiates these cases is part of the motivation behind the multi-machine, multi-discharge validation study presented in Chapter 6.

Now that all three validation comparisons have been made (ion heat flux, electron heat flux, and perturbative diffusivity) it is possible to reduce the overall agreement between simulation and experiment to a single parameter, also known as a validation metric. In particular this thesis will make use of the Ricci metric, χ_{Ricci} , originally described in Reference [149] and further explained in Appendix D. This χ_{Ricci} is not to be confused with the perturbative diffusivity χ_e^{pert} . The Ricci metric is essentially a weighted sum of the differences between the experimental and simulated quantities,

bound between 0 and 1.0. Lower values of χ_{Ricci} represent better agreement, and one can select a value below which the simulation is said to agree overall, depending on how strictly one wishes to define agreement. A cutoff somewhere between 0.1 and 0.5 is common. Using the three constraints in this validation give $\chi_{Ricci} = 0.05$, indicating excellent agreement. This agreement is better than the agreement calculated in the next validation study, though different constraints are used there.

Effect of Collisionality on Simulation Perturbative Diffusivity

Finally, the GENE simulations prepared for this validation study can also be used to investigate the experimentally observed trend of perturbative diffusivity with collisionality that was described in Chapter 3. Recall that experimentally, χ_e^{pert} was shown to scale as $(\nu_{GENE})^{-0.44}$ across both Alcator C-Mod and ASDEX Upgrade plasmas. The discharge in this validation study, ASDEX Upgrade 33585, was one of those used in that correlation, and is circled in green in Figure 3-9.

To try to understand why ion-scale gyrokinetic simulations agreed well with perturbative diffusivity in the study here but not in those earlier in this chapter, the GENE collisionality for the flux-matched simulation was artificially raised by an order of magnitude and the simulation was rerun with all other parameters the same. The input a/L_{Te} into this higher collisionality simulation was also scanned up and down 12%, and a perturbative diffusivity was calculated by a fit to set of three simulations.

This calculation resulted in a perturbative diffusivity of 0.23 m²/s, which is a factor of 38 lower than the nominal collisionality for the flux matched case, which gave 8.8 m²/s. Based on the experimental trend, $(\nu_{GENE})^{-0.44}$, one would expect the perturbative diffusivity to decrease by only a factor of roughly 3 when the collisionality is increased by a factor of 10 ($10^{-0.44} \approx 0.36$). The simulation therefore overestimates the strength of the dependence of the perturbative thermal diffusivity

on collisionality, predicting a much stronger decrease of perturbative diffusivity with increasing collisionality than is observed experimentally.

Note that this particular investigation of the dependence of the perturbative diffusivity on collisionality is not intended as a direct validation study, but instead as a search for insight into reasons for the dependence on collisionality. While it is possible to scan the experimental collisionality while keeping other dimensionless parameters constant, this particular study did not do so, instead just scanning the simulation collisionality without any reference to a specific experimental discharge. In the future, it may be worthwhile to perform such collisionality scans in order to investigate the dependence more directly.

One interesting result of the simulation collisionality scan is that the very small value predicted for the perturbative diffusivity, $0.23 \text{ m}^2/\text{s}$, is very similar to the values predicted by the ion-scale gyrokinetic simulations of higher collisionality Alcator C-Mod plasmas that were presented in the first half of this chapter. In other words, when the collisionality of the discharge considered in this validation study is increased to be similar to that observed in Alcator C-Mod, then a similar order of magnitude under-prediction of the experimentally expected perturbative diffusivity is observed. This suggests that perhaps the low collisionality of the plasma simulated in this work is related to the agreement found between ion-scale simulations and experiment, as opposed to the high collisionality plasmas considered in previous studies.

As will be discussed further at the end of this chapter, this may be related to the impact of the collisionality on the relative strength of ITG and TEM turbulence, as well as the balance between ion- and electron-scale phenomena. In other words, changing the collisionality changes the relative importance of different turbulent modes, impacting the ability of ion-scale simulations to accurately model the plasma behavior.

Conclusions and Discussion of Inner Radius GENE Validation on ASDEX Upgrade

To summarize the results of this particular validation study, local ($\rho_{tor} = 0.49$), non-linear, ion-scale GENE simulations were able to simultaneously match experimentally measured ion heat flux, electron heat flux, and electron perturbative thermal diffusivity within experimental uncertainty. Linear stability analysis revealed that this plasma is strongly mixed-mode in the ion-scale range, with ion- and electron-modes alternating as the dominantly unstable mode. In order to simultaneously match the ion and electron heat fluxes with nonlinear simulations, both a/L_{Ti} and a/L_{Te} were adjusted within experimental uncertainty. This same heat flux matched simulation also agreed with the experimentally measured perturbative diffusivity, measured experimentally using partial sawtooth heat pulses and calculated from the simulation by scanning a/L_{Te} . A scan of the simulation collisionality suggested that the low collisionality in this particular discharge may be related to the better match between simulation and experiment (likely due to its effect on the relative importance of different turbulent modes). This set of simulations is the first ion-scale simulation that has been able to match the experimental perturbative diffusivity.

Consider now how this result fits into the larger validation picture. As has been made increasingly clear by the validation studies presented in this chapter, and others in recent years, such as those in References [10, 11, 12, 13, 14, 15, 16, 17, 55, 64, 144], multi-scale effects and cross-scale coupling are in some instances absolutely vital in order to find agreement between gyrokinetic simulations and experiment. In other cases, however, such as the one presented here, ion-scale simulations are able to match the experiment without the inclusion of multi-scale effects. Cases have now been identified in which: ion-scale simulations miss both the experimental heat fluxes and perturbative diffusivity [14, 55, 100], and only multi-scale simulations are able to match the experiment [14]; ion-scale simulations match heat fluxes, but miss the

perturbative diffusivity, while multi-scale simulations match everything [17]; and ion-scale simulations simultaneously match heat fluxes and perturbative diffusivity (the study just presented).

In practical terms, this range of results indicates that it is not always necessary to run enormously expensive multi-scale simulations, and that instead ion-scale simulations are fully sufficient. The difficult part is determining when this is the case, without first running all of the simulations and comparing to experiment, as would be necessary when eventually using the simulations predictively for a future machine. This question is precisely the motivation for the study presented in Chapter 6, and as such will be discussed further at that point. The topic is, however, worth a brief discussion at this point, as the methodology presented later builds off of what was learned through the validation studies here.

Earlier in this validation study it was mentioned that the collisionality, and thus the balance of ITG and TEM turbulent modes at long wavelengths, may have some correlation to the ability of ion-scale simulations to match experimental measurements. This is investigated further in Chapter 6, and is shown to likely be one piece of a larger picture concerning the relative strength of different turbulent modes. Similarly, past work has suggested that the ratio of turbulence linear growth rates at high-wavenumber to those at low-wavenumber may also be related to the ability of ion-scale simulations to match experimental measurements [14, 17]. Intuitively, increasing the drive of electron scale turbulence relative to ion-scale turbulence would increase the importance of electron-scale phenomena, and thus the importance of multi-scale physics.

Specifically, one is interested in the ratio $\gamma_{high-k}/\gamma_{low-k}$, where γ_{high-k} is the maximum linear growth rate of any mode above $k_y\rho_s \gtrsim 2.0$, and γ_{low-k} is the maximum linear growth rate of any mode below $k_y\rho_s \lesssim 2.0$. This particular metric won't distinguish between the importance of purely electron-scale phenomena and the importance

of cross-scale coupling, as it just compares the electron- and ion-scale turbulent drive, but it should give some general indication of the importance of electron-scale turbulence. This distinction is addressed further in Chapter 6.

Reference [14] suggested that plasmas in which $\gamma_{high-k}/\gamma_{low-k} \lesssim 40$ may not have significant multi-scale effects, and that ion-scale simulations would therefore be sufficient. More recently, however, Reference [17] presented a case in which $\gamma_{high-k}/\gamma_{low-k} \approx 30$ in which multi-scale interactions were required in order to match the experimental perturbative diffusivity (but not the heat fluxes). The exact value of $\gamma_{high-k}/\gamma_{low-k}$ below which multi-scale effects are unimportant may therefore be somewhat less than the value of 40 originally proposed.

Using the linear stability results for the discharge presented in this validation study, one calculates $\gamma_{high-k}/\gamma_{low-k} \approx 19$ for the nominal experimental parameters and $\gamma_{high-k}/\gamma_{low-k} \approx 15$ for the flux matched simulation. Both of these values are much lower than for the discharges presented in past work, and far below the values of 30 or 40 originally suggested as cutoffs. The fact that ion-scale simulations agree in this case, but not in other with larger values of $\gamma_{high-k}/\gamma_{low-k}$ is consistent with the hypothesis that there is some cutoff value. Chapter 6 will expand upon this analysis, showing that a slightly modified criterion is better able to predict the importance of multi-scale effects. All of that analysis, however, was motivated by and expands upon the work presented in this chapter.

This completes the discussion of the GENE validation study at the inner-radius of ASDEX Upgrade discharge 33585. This chapter will present one final gyrokinetic validation study before moving on to some overall results and discussion of what was learned through this work.

5.2.3 Outer Radius Validation Study with CECE Measurements

The second validation study presented in this thesis using GENE on ASDEX Upgrade focused on $\rho_{tor} = 0.75$. This study was not performed directly as part of this thesis, but included significant contributions by the author of this thesis and utilized both hardware and analysis from this thesis. For this reason, somewhat less detail is given for this study than for the first validation study in this section. The work was originally presented in Reference [107].

At this radial location, measurements of heat fluxes (from TRANSP), electron temperature fluctuations (from CECE), radial correlation lengths (from CECE), and n-T phase (from CECE and reflectometry) were available. While this thesis is not generally focused on the details of radial correlation length and n-T phase measurements, a brief description of both is given here.

Radial Correlation Length and n-T Phase Measurements

Radial correlation length measurements are based on all of the same fundamental principles as normal CECE measurements, but instead of just two radially separated channels, they require a set of many radially separated channels. In traditional CECE measurements, one wishes to have the two radiometer channels as close as possible without actually overlapping, in order to ensure that one is measuring the same turbulent features without measuring the same thermal noise. In order to make radial correlation length measurements, one uses a series of many radially spaced channels, and calculates the coherence between the first channel and the second, the first and the third, the first and the fourth, and so on. In this manner, one is able to measure how the correlation between channels decreases as the channel spacing increases. The fall-off of this coherence is a measure of the radial correlation length.

Since the upgraded CECE system on ASDEX Upgrade contains a comb of many closely spaced channels, no hardware adjustments are necessary in order to measure radial correlation lengths. One simply uses the same data that was recorded for the CECE measurements in order to also calculate the radial correlation length. This is one of the key strengths of such a system design.

In addition, one can use a CECE system in concert with a reflectometer to measure the phase angle between temperature and density fluctuations at the same radial location. Both theoretical and experimental work indicates that the n-T phase angle is intimately linked to the fundamental nature of different turbulent modes [150, 151, 152]. Thus, by measuring the n-T phase angle, one is able to constrain the relative mix of ion temperature gradient (ITG) and trapped electron modes (TEM) at long wavelength. This provides yet another validation constraint on gyrokinetic simulations, as one cannot arbitrarily change the mix of the long wavelength turbulence in order to match the heat fluxes or other constraints.

Experimentally, n-T phase measurements are relatively difficult and have only been published in a few select instances [150, 151, 152]. While there are quite a few intricacies of the measurement that will not be covered here, the general idea is that one measures density fluctuations and temperature fluctuations at the same location at the same time and then calculates the phase angle between the two sets of data. This is done by coupling a CECE radiometer and a reflectometer through the same optics. The CECE system passively collects ECE from the plasma, while the reflectometer actively launches a wave into the plasma, which is then reflected at the wave's cutoff frequency. The location of the reflection depends on the type of wave launched (X- or O-mode) and the density of the plasma. In particular, the CECE system on ASDEX Upgrade is coupled to the reflectometer systems described in References [106, 153, 154].

One of the main experimental difficulties, then, is ensuring that the reflectometer

cutoff aligns exactly with one of the CECE measurement locations, especially considering uncertainties in the density measurements in tokamak plasmas. Again, this task is made easier by the large number of channels in the ASDEX Upgrade CECE system, which cover a large radial region of the plasma, as the reflectometer can align with any of the channels in the system. The resulting cross-phase calculation, using the phase of the cross-correlation between the reflectometer and radiometer signals, gives a phase angle spectrum, as the n-T phase angle may depend on the fluctuation frequency, which can then be compared to the outputs of gyrokinetic simulations.

Outer Radius GENE Validation Results on ASDEX Upgrade

As discussed above, most of the GENE simulation settings were exactly the same in this study as in the inner-radius study, with the exception of the local plasma parameters. These parameters are summarized in Table 5.4. As stated earlier, the GENE simulations in this study were not performed directly as part of this thesis, so this section is somewhat more condensed than the last.

While many simulations were run in order to best match the validation constraints, in the interest of brevity, only the results from the best match to the constraints will be presented. As in the inner-radius study presented above, both the ion and electron temperature gradients were varied within experimental uncertainty in order to match the validation constraints. Overall, the best agreement was found for the nominal value of a/L_{Ti} (+0%) and a/L_{Te} reduced by 13% (which is within the experimental uncertainty of 20%). The five validation constraints compared here are the electron and ion heat fluxes, the electron temperature fluctuations, the electron temperature fluctuation radial correlation length, and the density-temperature fluctuation phase angle.

As was discussed for the inner-radius study, heat fluxes are obtained by averaging the outputs of the GENE simulation over many a/c_s times in order to ensure that

an accurate mean value is found. With the inputs described above (nominal a/L_{Ti} and a/L_{Te} reduced by 13%), the simulation overpredicts the ion heat flux by 13% (compared to an uncertainty of 26%), and under-predicts the electron heat flux by 14% (compared to an uncertainty of 18%)³. These are both within experimental uncertainty.

Consider next the comparison with the measured temperature fluctuation spectrum. As in the study with GYRO on Alcator C-Mod, a synthetic diagnostic must be applied to the outputs of GENE in order to meaningfully compare with experimental measurements. As in that case, a two-dimensional Gaussian emission volume was chosen to model the diagnostic measurement region. An ECE radiation transport model [155] was used to estimate the radial dimension, giving a $1/e$ radial width of 4.7 mm. This model takes into account the finite filter bandwidth, relativistic broadening, Doppler broadening, and beam refraction.

The beam width (the poloidal extent of the emission volume) was determined via ex-situ measurements of the system optics, described in greater detail in [107]. These measurements were made with a 100 GHz source, and then extrapolated to the system frequency of 117 GHz using standard quasi-optical beam propagation techniques. These measurements gave a $1/e$ electric field beam radius of approximately 15 mm at the measurement location of interest in this plasma.

A final consideration in this study on ASDEX Upgrade that was not pertinent to the study on Alcator C-Mod regards the differences between perpendicular and parallel temperature fluctuations. Since the magnetic field in a tokamak introduces anisotropy along the magnetic field, the electron temperature fluctuations parallel and perpendicular to the magnetic field are not necessarily equal. This effect is more pronounced in less collisional plasmas, in which the plasma takes longer to thermally

³Note that the uncertainty in P_e quoted here (18 %) is larger than that in Reference [107] (14 %). This is due to additional uncertainty discovered in the routines that calculate the radiated power profiles since the publication of that study. This increase in uncertainty only strengthens the statements made in Reference [107].

equilibrate, and is thus more relevant on ASDEX Upgrade than on Alcator C-Mod. Since CECE is sensitive to the perpendicular temperature fluctuations [73], one should really only be comparing the perpendicular fluctuations predicted by the simulation to the experimental measurements.

GENE has the capability to separately resolve the parallel and perpendicular fluctuations, making it ideal for comparison to CECE on a lower collisionality tokamak. In the particular plasma under consideration, the perpendicular temperature fluctuation level was 33% larger than the total temperature fluctuation level, making clear the necessity of such a distinction. For this reason, the comparison made in this section will utilize the perpendicular fluctuation level, and not the total fluctuation level. See Reference [139] for another study that discusses the differences between perpendicular and parallel temperature fluctuations in GENE. Note also that the distinction between perpendicular and total temperature fluctuations will be important in Chapter 6, for which an approximate correction will be made for the ASDEX Upgrade discharges (since the TGLF code, used in that study, is unable to separately resolve parallel and perpendicular fluctuation levels).

With all of these factors taken into account, the final comparison between the measured and simulated temperature fluctuation spectra are shown in Figure 5-13. As before, uncertainty is represented as the shaded regions. Note that this particular figure is shown as cross-power per unit frequency plotted against frequency, which tends to exaggerate differences between two traces (since it is proportional to the temperature fluctuation level squared). This figure clearly shows that the simulations over-predict the measured fluctuation level, as well as over-predicting the spectrum's peak frequency. More quantitatively, when one integrates the total fluctuation level using the techniques described in Chapter 4, one finds a measured fluctuation level of $(\tilde{T}_{e,\perp}/T_e)_{exp} = 0.70 \%$ and a simulation value of $(\tilde{T}_{e,\perp}/T_e)_{GENE} = 1.14 \%$. The simulation therefore over-predicts the integrated fluctuation level by roughly 60%, far




Figure 5-13: Comparison of experimental and simulated electron temperature fluctuation spectra from $\rho_{tor} = 0.75$ in ASDEX Upgrade discharge 33585. Experimental measurement is in turquoise and GENE output is in orange. Only perpendicular temperature fluctuations are considered. Uncertainty is represented by the shaded regions. Data integrated from 2 seconds of plasma operation. Figure modified from Reference [107].

outside of the experimental uncertainty of approximately 2%.

Consider next the radial correlation length comparison, shown in Figure 5-14. As described above, the experimental measurement is accomplished using the comb of CECE channels, and correlating one channel with each subsequent channel in turn. This figure is plotted as the cross-correlation (in arbitrary units) against the channel separation in mm. The cross-correlation is similar to the complex coherence described in Chapter 4, but is not normalized since only the spatial variation of the signal is important for the correlation length measurement. Fitting a Gaussian to the results shown gives $L_r(T_{e,\perp})_{exp} = 9.8$ mm and $L_r(T_{e,\perp})_{GENE} = 10.5$ mm. For reference, 10 mm is approximately $6 \rho_s$ in this plasma. The uncertainty in this measurement is approximately 1.5 mm, so in this case the experimental and simulation results agree within uncertainty. Further discussion of whether or not a Gaussian is an appropriate shape for this correlation can be found in Reference [107].




Figure 5-14: Comparison of experimental and simulated electron temperature fluctuation radial correlation length measurements for ASDEX Upgrade discharge 34626 (a repeat of 33585). Experimental measurement is in turquoise and GENE output is in orange. Figure modified from Reference [107].

Finally, consider the phase angle between electron density and temperature fluctuations, for which Figure 5-15 shows the comparison. Again, this measurement couples the CECE system described in this thesis with reflectometers described in References [106, 153, 154]. In general, one should only compare the measured phase angle in the region in which there is strong coherence between the CECE and reflectometer signals; in this case, between approximately 10 and 50 kHz. As Figure 5-15 shows, GENE agrees quantitatively with the experimental measurement of a phase angle of approximately $\alpha_{nT} \approx -90^\circ$ over this frequency range.

As described above, the density-temperature phase angle is a good means by which to distinguish different turbulent modes (ITG and TEM) at long wavelength. For reference, GENE simulations of this discharge for which $a/L_{Te} = 0$ (ITG-like) gives $\alpha_{nT} = -135^\circ$. A simulation in which $a/L_{Ti} = 0$ (TEM-like) gives $\alpha_{nT} = -15^\circ$. The experimental value of $\alpha_{nT} \approx -90^\circ$ thus represents a value in between pure ITG and pure TEM (as one might expect) and the agreement with the GENE result indicates that GENE is properly capturing this mix.




Figure 5-15: Comparison of experimental and simulated density-temperature phase angle from ASDEX Upgrade discharge 34626 (a repeat of 33585). Experimental measurement is in turquoise and GENE output in orange. Uncertainty is represented by error bars. Figure modified from Reference [107].

To combine these validation constraints into a single metric, this study will again use the Ricci metric, χ_{Ricci} , that was introduced in the last study. For the simulation described here (with the nominal a/L_{Ti} and a/L_{Te} reduced by 13%), the five validation constraints (electron heat flux, ion heat flux, electron temperature fluctuation level and radial correlation length, and n-T phase angle) together give a Ricci metric of $\chi_{Ricci} = 0.27$. While this is a weaker agreement (a larger value of χ_{Ricci}) than for the inner-radius study presented above, it is still below 0.5, and thus should still be interpreted as an overall agreement.

To summarize this outer-radius validation study, a local, nonlinear, ion-scale GENE simulation at $\rho_{tor} = 0.75$ in ASDEX Upgrade discharge 33585 was able to simultaneously match the experimental electron heat flux, ion heat flux, electron temperature fluctuation radial correlation length, and density-temperature fluctuation phase angle within experimental uncertainty. The simulation did disagree with the measured total fluctuation level, but given the good agreement with the other four constraints, a combined validation metric, χ_{Ricci} , indicates that overall, the sim-

ulations agreed with the experiment.

There are a number of possible explanations for the discrepancy in the electron temperature fluctuation level from experiment and simulation. One possible issue is the CECE beam size, which can have a large impact on the total measured fluctuation level, as was described in Chapter 4. Reference [107] describes in detail the measurements of the beam size, and also reports that the beam size would need to be a factor of 2 smaller than the measured value in order to agree with the GENE results, which is far outside of the uncertainty in the beam measurements. Another explanation is that since the ion-scale GENE simulations do not include electron-scale heat transport, the simulations are over-driving low-wavenumber turbulence in order to match the heat fluxes, and in doing so are over-predicting the low-wavenumber fluctuation level. Finally, it is possible that other experimental parameters, such as the effective charge and density gradient, are impacting the results of the simulations, and that many experimental parameters would have to be changed simultaneously in order to match the experimental conditions. Chapter 6 will attempt to address the second and third possibilities through the use of multi-scale TGLF and multi-parameter optimization of the code outputs.

5.3 General Validation Observations and Conclusions

Taken together, the results of the four validation studies presented in this chapter lead to several overall conclusions, and motivate the work presented in Chapter 6. The results of the four validation studies are first collected and summarized here, before then considering how they fit into a larger picture.

Of the four validation studies in this chapter, two were with GYRO on Alcator C-Mod and two were with GENE on ASDEX Upgrade. All four studies employed electron and ion heat fluxes as validation constraints, and each included at least one other constraint. Most of the plasmas were L-modes, though two studies included

both L- and I-mode plasmas. The first study with GYRO on Alcator C-Mod included an L- and I-mode plasma, showing that global ion-scale GYRO simulations matched the ion heat flux for both plasmas, but missed the electron heat flux, perturbative diffusivity, and electron temperature fluctuations in both plasmas (though the I-mode electron heat flux was matched in some regions, and the temperature fluctuation level was below the instrumental sensitivity). The second study involved adding the perturbative thermal diffusivity to two previously completed sets of simulations, showing that in at least one case, including multi-scale effects was able to resolve the discrepancy between ion scale simulations and experimentally measured electron heat flux and perturbative diffusivity. In addition, this study showed that even in cases where ion-scale simulations match experimental heat fluxes, they may still disagree with the experimental perturbative diffusivity.

The third study, with GENE on ASDEX Upgrade, showed that an ion-scale simulation was able to successfully match electron and ion heat fluxes and the perturbative diffusivity at an inner radius of one discharge. The final study showed that ion-scale GENE simulations successfully matched electron and ion heat fluxes, electron temperature fluctuation correlation lengths, and density-temperature fluctuation phase angle, but missed the electron temperature fluctuation level at an outer radius in the same ASDEX Upgrade discharge.

The most basic conclusion from these studies is that ion-scale gyrokinetic simulations are clearly not always sufficient to describe plasma behavior in tokamaks, though in some cases they are. The question of why the simulations disagree with experiment is not always particularly straightforward, as there could be many possible reasons. One of the central motivations behind these studies, however, was the relatively recent work with multi-scale simulations and indications that in some plasma conditions, multi-scale simulations agree with validation constraints while ion-scale simulations disagree. For the second study presented here, it was directly shown that

a multi-scale simulation agrees with the experimental constraints, while an ion-scale simulation disagrees. Though no multi-scale simulations were performed for the other three studies, these three studies show that ion-scale simulations are sometimes sufficient to model the plasma behavior, but are in other cases insufficient. Specifically, experimental cases have now been identified in which: ion-scale simulations disagree with both heat fluxes and perturbative diffusivity, ion-scale simulations agree with heat fluxes but disagree with perturbative diffusivity, and ion-scale simulations agree with both.

Another result of this chapter is that the perturbative diffusivity seems to be a powerful validation constraint. In addition to its use in three of the four validation studies presented here, in Reference [17] the perturbative diffusivity was the only parameter that distinguished between ion- and multi-scale simulations. Recall that the perturbative diffusivity correlates experimentally with the plasma collisionality, and that this correlation is possibly connected to the balance between ITG and TEM turbulent modes.

Digging a little deeper, it is possible to interpret the perturbative diffusivity as a measure of the balance between ITG, TEM, and ETG modes in the plasma. The electron perturbative thermal diffusivity is measured through the propagation of *electron* heat pulses, which dominantly change the electron temperature gradient. The value of the perturbative thermal diffusivity should therefore be greater when the dominant turbulence is sensitive to the electron temperature gradient, such as when TEM or ETG modes are active. If the dominant turbulence is ITG, the perturbative diffusivity should be lower, as the heat flux should be less sensitive to the electron temperature gradient. In this sense, the perturbative diffusivity constrains the relative mix of ITG and TEM/ETG driven heat fluxes in the simulation.

With this interpretation, it is unsurprising that an ion-scale simulation dominated by ITG turbulence would severely under-predict the perturbative diffusivity, as it is

unable to include the contributions of the ETG. The plasma analyzed in the third validation study, however, was dominantly TEM at low wavenumbers for the flux matched inputs, which may be responsible for the better agreement between simulation and experiment. In particularly mixed mode plasmas, cross-scale coupling may become important, so it may be important to include multi-scale effects in order to correctly discern the perturbative diffusivity. While the plasma from the third validation study was to some extent mixed mode, at the heat flux matched inputs, the plasma was dominantly TEM and ETG at all wavenumbers, so it may be that the cross-scale coupling was unimportant.

Turning back to the larger validation picture, the results presented here raise the question of *why* ion-scale simulations agree with experiment in some cases and disagree in others. In other words, what properties of the plasmas in question distinguish cases in which ion-scale simulations are sufficient from those in which they are not? One possibility that was suggested during the third validation study (inner-radius ASDEX Upgrade with GENE), was that the ratio of high-k to low-k peak linear growth rates, $\gamma_{high-k}/\gamma_{low-k}$, may be a distinguishing parameter. The idea is that at higher values of $\gamma_{high-k}/\gamma_{low-k}$, electron-scale turbulence is more strongly driven compared to ion-scale turbulence, and thus multi-scale effects are more important and ion-scale simulations perform less well. That particular validation study had $\gamma_{high-k}/\gamma_{low-k} \approx 15$, much lower than the cutoff proposed in References [14, 17] of $\gamma_{high-k}/\gamma_{low-k} \approx 30-40$. While this data was not available for the two validation studies not performed directly as part of this thesis (the first and fourth), the second study (GYRO on L- and I-mode Alcator C-Mod plasmas) did have this data. For the L-mode discharge, where the ion-scale GYRO simulations strongly disagreed with the experiment, $\gamma_{high-k}/\gamma_{low-k} \approx 12$, which does *not* fit into the proposed model. In that case, a plasma with very low $\gamma_{high-k}/\gamma_{low-k}$ still required multi-scale effects to find agreement.

Though this is a very limited data set, it is clear that while $\gamma_{high-k}/\gamma_{low-k}$ may

have some correlation to the importance of multi-scale effects, the real picture must be more complicated. In order to investigate this question, one would need a larger set of validation discharges, preferably from multiple machines and using the same code, to avoid the question of machine or code dependence. Obtaining such a large data set using GENE or GYRO would, however, take an enormous amount of computing resources. This is exactly the motivation for the study presented in Chapter 6, which uses TGLF to validate 11 plasma discharges, some at two locations for a total of 17 plasma conditions, from Alcator C-Mod and ASDEX Upgrade. The large data set gathered in that study expands up on the results of this chapter, revealing a set of criteria with which one can predict whether or not multi-scale effects will be important in a plasma discharge, based only on linear stability results.

Chapter 6

Validation of TGLF in Multi-Machine Study

This chapter seeks to answer more concretely one of the central questions of this thesis, concerning when multi-scale effects are important in turbulent transport simulations and when ion-scale simulations are sufficient. The work presented in this chapter will build off of the work of the last three chapters. It will make use of both perturbative diffusivity (from Chapter 3) and CECE (from Chapter 4) as validation constraints. The validation of gyrokinetic simulations in Chapter 5 motivated this work from from a physics standpoint, showing that ion-scale simulations can in some cases agree with many experimental validation constraints, while in others ion-scale simulations are entirely insufficient. That work also motivated the methodology of this chapter, as multi-scale gyrokinetic simulations are far too computationally expensive to perform on more than a few discharges.

This chapter makes use of the quasi-linear turbulent transport code TGLF (Trapped Gyro-Landau Fluid) [20, 21], which was described in greater detail in Chapter 2. For the highest fidelity predictions of future machine performance, especially in parameter regimes far from current plasmas, fully nonlinear gyrokinetic simulations are likely

required, but TGLF is used here as a tool to quickly assess the importance of multi-scale effects (though from a practical perspective, it is always better to run TGLF in the multi-scale configuration, as the increase in computational time is negligible). The chapter begins by applying TGLF to the L- and I-mode validation case that was originally studied with GYRO presented in Section 5.1.1. Using the results of this application as motivation, the chapter then presents a full validation study of TGLF on 11 discharges from Alcator C-Mod and ASDEX Upgrade.

This validation study reveals first that TGLF run in a multi-scale configuration agrees with all available validation constraints on all of the experimental plasma conditions under consideration. Going further, the code is also run in an ion-scale configuration on each condition, and a comparison of the two configurations leads to a set of criteria that determine when multi-scale effects are necessary for agreement with experimental constraints. In addition to their practical implications, these criteria reveal which physical processes are dominantly responsible for the importance of multi-scale effects. Finally, this chapter uses both ion- and multi-scale TGLF to predict profiles for two experimental discharges, in order to illustrate the practical implications of validation when one's ultimate goal is the prediction of future machine performance.

6.1 TGLF Applied to L- and I-mode Validation Study

Before embarking on an entirely new validation study with new experimental conditions, TGLF is first used to expand upon the validation effort originally presented in Section 5.1.1. Recall that this validation effort originally ran GYRO on an L-mode time period and an I-mode time period from the same Alcator C-Mod discharge, 1120921008. For convenience, the overall plasma parameters are repeated here: $\bar{n}_e = 0.7 - 0.8 \times 10^{20} \text{m}^{-3}$, $I_p = 1.1 \text{ MA}$, $B_t = 5.4 \text{ T}$, and $q_{95} = 3.2$. The plasma was heated with 1.6 MW of ICRH in the L-mode time period and 3.5 MW of ICRH


in the I-mode time period. As with the GYRO work, this study will compare experimental ion and electron heat fluxes and electron perturbative thermal diffusivity to the outputs of TGLF. It will not, however, compare electron temperature fluctuations, though this constraint will be included in the multi-machine validation study later in this chapter. The results of this work with TGLF were originally published in Reference [100].

The original motivation of this work was to see whether the multi-scale TGLF model would agree with the experimental validation constraints when ion-scale GYRO was unable to do so. While it is recognized that such a comparison is not strictly rigorous, as one is comparing two different simulation codes, it was intended as a quick test of the hypothesis that multi-scale effects would resolve the discrepancies observed in the ion-scale GYRO results. The results of this check further informed the methodology that was used in the more extensive validation study presented later in this chapter.


In this study, TGLF was run using its standard multi-scale wavenumber spectrum ($k_{\theta}\rho_s = 0.1$ to 24.0) and the SAT-1 saturation rule [21]. Ion and electron heat fluxes are calculated as standard outputs from TGLF, and so can be compared to the experimental TRANSP analysis directly. The ion and electron temperature gradients were varied within experimental uncertainty in order to match the heat fluxes as closely as possible. As was done for the gyrokinetic codes in Chapter 5, the perturbative diffusivity is calculated by running TGLF cases with the input electron temperature gradient increased and decreased by 12%, and then calculating the slope of the output electron heat flux against the input electron temperature gradient.

Heat Flux Comparison


Unlike GYRO, TGLF cannot be run as a global simulation, so instead of comparing global results, TGLF was run at three radii ($\rho_{tor} = 0.64, 0.74, 0.84$) for both L- and



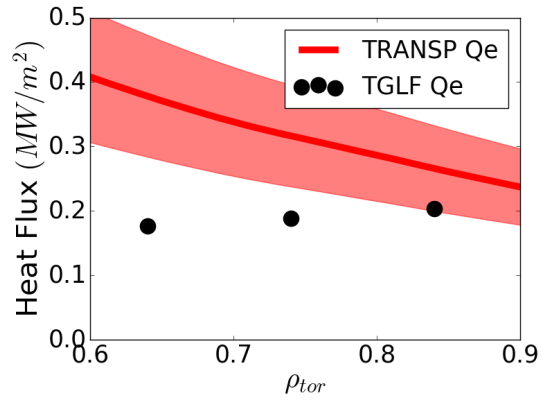
(a) L-mode ions.



(b) L-mode electrons.



(c) I-mode ions.



(d) I-mode electrons.

Figure 6-1: Comparison of the experimental and TGLF electron and ion heat fluxes in the L- and I-mode periods. Ion heat fluxes are shown on the left, and electron heat fluxes on the right. L-mode results are shown on the top in blue, and I-mode on the bottom in red. TGLF results are shown as the black circles. TGLF generally matches ion heat fluxes (with one slight deviation in I-mode), but significantly underpredicts the electron heat fluxes in both cases. Figure originally from Reference [100].

I-mode experimental cases. Figure 6-1 shows the comparison between TGLF and the experimental electron and ion heat fluxes from both L- and I-mode time periods of the discharge. These results are not substantially different from the GYRO validation results shown in Figure 5-4. Similar to the GYRO outputs, TGLF is able to match the ion heat flux in both L- and I-mode within experimental uncertainty, except for a slight disagreement at the outermost radius in I-mode. For the electron heat flux, TGLF is much closer than GYRO for the L-mode case, but still slightly underpredicts the electron heat flux at all radii. In I-mode, TGLF only matches the experimental electron heat flux at the outermost radius, the same as for the GYRO results.

This disagreement contrasts with previous results using multi-scale TGLF SAT-1, in which multi-scale TGLF was able to reproduce both experimental heat fluxes, Q_i and Q_e , in an Alcator C-Mod L-mode plasma for which ion-scale simulations were insufficient [21]. This disagreement suggests that there may be additional effects, beyond just the inclusion of multi-scale effects, that lead to disagreement between TGLF (or other turbulent transport models) and experimental measurements. Specific examples are given at the end of this section, and inform the validation methodology of the multi-machine study presented later in this chapter.

Perturbative Thermal Diffusivity Comparison

In addition to electron and ion heat fluxes, the perturbative thermal diffusivity was used as a validation constraint for these TGLF simulations. The perturbative diffusivity calculated from the simulation electron temperature gradient scans are given in Table 6.1. For convenience, this table also includes the perturbative diffusivities calculated with GYRO that were shown in Section 5.1.1. The TGLF perturbative diffusivity shown in this table is an average of the values calculated at each of the three radial locations, in order to most accurately compare to the experimental measurement.

	L-mode	I-mode
	$\chi_e^{pert} (m^2/s)$	$\chi_e^{pert} (m^2/s)$
Experiment	4.0 ± 0.6	4.3 ± 0.9
GYRO	0.4	1.0
TGLF	0.7	0.6

Table 6.1: Comparison of perturbative diffusivity from experiment χ_{Exp}^{pert} , GYRO χ_{GYRO}^{pert} , and TGLF χ_{TGLF}^{pert} in both L- and I-mode. All values given in units of m^2/s and are average of the values calculated at each of the three radial locations. Originally from Reference [100].

As was the case with ion-scale GYRO, multi-scale TGLF severely underpredicts the perturbative diffusivity in both L- and I-mode. While the multi-scale TGLF prediction is slightly closer in L-mode, it is further away in I-mode. This stark disagreement makes clear that simply including multi-scale effects is not always sufficient to attain agreement with experimental validation constraints, and that there are clearly other effects of importance in these validation efforts.

Conclusions from Preliminary Multi-Scale TGLF Validation

At first glance, the results of this preliminary work with multi-scale TGLF, applied to the same L- and I-mode plasmas investigated with GYRO in Section 5.1.1, would seem to indicate that the hypothesis that multi-scale effects are responsible for the disagreement between many of the gyrokinetic results in Chapter 5 and experiment is false. This seems to at least be the case in the particular plasmas considered here. The multi-machine study presented later in this chapter seeks to answer this question more directly, by comparing ion- and multi-scale TGLF run on the same plasma discharges using the same validation constraints.

The results from this preliminary validation, however, motivate another possible explanation for some of the discrepancy observed in previous validation efforts. As

has been mentioned before in this thesis, the computational resources required to run nonlinear gyrokinetic simulations, even ion-scale simulations, typically limit one to varying only the input temperature gradients, and often just the input ion temperature gradient, within experimental uncertainty in order to try to match the heat fluxes. Only very few studies also vary other inputs, such as the density gradient, effective charge, T_e/T_i , collisionality, and others. The very fast runtime of TGLF, however, makes such multi-parameter scans feasible. For this reason, in addition to comparing the results of ion- and multi-scale TGLF, the multi-machine validation study presented next will also compare results of varying only temperature gradient inputs or of varying temperature gradients, density gradient, and effective charge, Z_{eff} . To foreshadow the results of those scans, both multi-scale effects and the variation of input parameters beyond just temperature gradients prove necessary to match experimental constraints in some cases.

6.2 Methodology and Plasmas for Multi-Machine Study

The final validation study of this thesis differs significantly in both its objective and methodology from the previous studies. In terms of objective, this final study intends to do more than just validate a particular code on a particular plasma or set of plasmas. While this will indeed be accomplished along the way, the real objective is to develop a set of criteria with which one can determine when multi-scale effects are important to turbulent transport in tokamak plasmas, and when ion-scale simulations are sufficient. In this sense, this final validation study has a broader physics-oriented objective than the previous studies. The form of these criteria will also shed light on the underlying physical interactions that determine when multi-scale interactions are important.

To accommodate this broader objective, the methodology of this final study also differs significantly from past validation work. Instead of focusing in depth on a

single discharge or pair of discharges, this study will look at eleven plasmas from two machines (Alcator C-Mod and ASDEX Upgrade), some at two radial locations for a total of 17 plasma conditions. As with previous studies, a total of four validation constraints will be applied: electron heat flux, ion heat flux, perturbative thermal diffusivity, and electron temperature fluctuations. Unlike previous studies, however, and motivated by the results of the gyrokinetic results from Chapter 5 and the TGLF work from Section 6.1, TGLF will be run in both ion- and multi-scale configurations to investigate the importance of multi-scale effects, and multiple inputs will be varied within uncertainties.

The remainder of this section describes the experimental plasmas considered in this study and gives further detail on the validation methodology.

Table 6.2 summarizes the 11 plasma discharges from Alcator C-Mod and ASDEX Upgrade that are analyzed in this study. These discharges are all low to moderate power L-modes, and were selected to avoid any nonstandard operating conditions, such as internal transport barriers, reversed magnetic shear, resonant magnetic perturbations (RMPs), etc. Part of the reason that only L-mode plasmas were analyzed in this study, instead of H- or I-modes, was to allow more consistent measurement of the electron temperature fluctuation level, as fluctuation levels in I- and H-modes are often below the instrumental sensitivity of CECE diagnostics. The Alcator C-Mod plasmas were all heated with hydrogen minority ion cyclotron resonance heating (ICRH) and the ASDEX Upgrade plasmas were all heated with electron cyclotron heating (ECH), with neutral beam blips for charge exchange measurements.

Availability of perturbative thermal diffusivity measurements (via partial sawtooth heat pulses) and electron temperature fluctuation measurements (via CECE) was also important to the selection of these plasmas. The first four Alcator C-Mod discharges have both measurements available at the same radial location. The last Alcator C-Mod plasma has only CECE measurements. All six ASDEX Upgrade plas-

Machine	Discharge	I_p (MA)	P_{aux} (MW)	Radius	n_e (10^{19} m^{-3})	T_e (keV)	\tilde{T}_e/T_e (%)	χ_e^{pert} (m^2/s)
Alcator C-Mod	1120706008	0.8	1.2	0.75	6.52	0.52	1.05	3.21
	1120706017	0.8	1.2	0.75	8.83	0.45	0.81	1.85
	1120706018	0.8	1.2	0.75	10.0	0.41	0.78	1.76
	1120706019	0.8	1.2	0.75	6.28	0.51	1.08	3.41
	1120706030	0.8	4.5	0.75	6.74	0.71	1.09	NA
ASDEX Upgrade	33585	1.0	0.7	0.75	1.41	0.58	0.74	NA
				0.49	1.99	1.48	NA	7.72
	34301	0.8	0.3	0.70	1.99	0.40	0.45	NA
				0.39	2.46	1.36	NA	6.58
	34303	0.6	0.3	0.70	1.58	0.31	0.75	NA
				0.30	2.70	1.49	NA	3.85
	34309	1.0	0.6	0.70	1.85	0.47	0.69	NA
				0.50	2.47	1.15	NA	7.35
	34508	0.6	0.5	0.70	1.80	0.31	0.66	NA
				0.30	2.91	1.62	NA	6.00
	34623	0.6	0.6	0.65	2.50	0.30	0.41	NA
				0.27	3.90	1.38	NA	5.05

Table 6.2: Summary of plasma parameters for the five Alcator C-Mod and six ASDEX Upgrade discharges used in this study. All of the Alcator C-Mod discharges operated at 5.4 T magnetic field on axis, and all of the ASDEX Upgrade discharges at 2.5 T. Modified from Reference [65].

mas have both measurements available, but in every case they are at different radial locations (the fluctuation measurements are always further out radially than the perturbative diffusivity measurements). Temperature fluctuation measurement positions are constrained by the minimum measurable fluctuation level with CECE, and perturbative diffusivity measurements are constrained radially by the sawtooth mixing radius and the pulse amplitude below which the pulse is indistinguishable from noise. For this reason, each ASDEX Upgrade plasma is validated at two separate radial locations, each with three validation constraints (ion and electron heat fluxes plus either fluctuations or perturbative diffusivity), for a total of 17 plasma conditions. As will be discussed further in the next section, validation constraints are added one by one in the simulation runs, so it is possible to compare only validation efforts that used exactly the same constraints. It is also possible to compare the results of each condition with all of the constraints available for that condition, which is perhaps a simpler comparison.

In addition to a larger number of plasma conditions, TGLF will be run many times for each condition using different input variations, constraints, and scale range. In order to optimize up to four inputs in order to match up to four validation constraints, TGLF is run through the VITALS framework [156]. VITALS utilizes surrogate based machine learning algorithms in order to efficiently optimize a variable number of inputs to a turbulent transport model (within their experimental uncertainty) in order to match validation constraints as closely as possible. VITALS has been implemented using the OMFIT code [157], which serves as an interface to many other codes and data sources.

As mentioned above, in order to investigate directly the importance of multi-scale effects in different plasmas, TGLF was run in ion-scale and multi-scale configurations for each experimental case considered. In this study, the multi-scale TGLF runs include modes from $k_{\theta}\rho_s = 0.1$ to 24.0 (the standard wavenumber spectrum in TGLF),

where k_θ is the poloidal wavenumber (defined as in GYRO) and ρ_s is the ion gyro-radius evaluated at the electron temperature. The ion-scale TGLF configuration ran with modes from $k_\theta \rho_s = 0.08$ to 1.0.

If one uses the standard wavenumber spectrum, the simulation modes are spaced logarithmically, such that there is higher resolution at lower wavenumber. This is done so as to better resolve both ion- and electron-scale dynamics, as the variation with wavenumber is greater at lower wavenumber than at higher wavenumber. The same strategy is employed for linear growth calculations with GYRO and GENE, as can be seen in Chapter 5 in Figures 5-3 and 5-8. On the other hand, if one manually specifies a wavenumber range in TGLF, as was done for the ion-scale configuration, one must use evenly distributed wavenumbers. Specifically, this work used 12 evenly spaced wavenumbers from 0.08 to 1.0, which approximates the number of wavenumbers in this range used in the multi-scale configuration.

Note that in both ion- and multi-scale configurations, TGLF was run with the ‘SAT-1’ saturation rule [21]. This means that the rule used to determine the saturated turbulence levels from the linear growth rates (since TGLF is a quasi-linear model) is based on the larger set of GYRO discharges, including some multi-scale runs, described in Reference [21], as opposed to the original set of discharges that made up SAT-0 in Reference [20].

In addition to running TGLF in either ion- or multi-scale configurations, based on the results in Section 6.1 the number of inputs that are allowed to vary within uncertainties will also be changed. First, only a/L_{Te} and a/L_{Ti} will be allowed to vary by VITALS, as these are traditionally the drive terms that are varied in most validation studies. VITALS will then run a second validation effort for each case, now allowing a/L_{Te} , a/L_{Ti} , a/L_{ne} and Z_{eff} to all vary within uncertainty. This hopes to address the issue discovered in Section 6.1, where even a multi-scale TGLF run was unable to match experimental validation constraints if one only varied a/L_{Te}

and a/L_{Ti} . These four inputs are chosen since they are independent of one another (unlike, for example, T_i/T_e and ν_c , which both depend on temperature).

Finally, for each validation case, the available validation constraints (Q_e , Q_i , \tilde{T}_e/T_e and χ_e^{pert}) will be added consecutively, in order to investigate the possibility of fortuitous agreement when one applies too few constraints. Each case will start with only Q_e and Q_i as constraints, and will be run through VITALS to see if agreement is possible. If only three validation constraints are available, the third constraint will then be added. If all four constraints are simultaneously available, validation will be performed with first just Q_e and Q_i ; then Q_e , Q_i , and \tilde{T}/T ; then Q_e , Q_i , and χ_e^{pert} ; and finally with all four of Q_e , Q_i , \tilde{T}/T and χ_e^{pert} .

Each of these validation constraints is extracted from TGLF in a slightly different manner. The ion and electron heat fluxes are a standard output from TGLF, and therefore don't require any special procedure [20, 21]. The perturbative diffusivity is calculated using the same procedure as was used for the gyrokinetic codes in Chapter 5: the input electron temperature gradient was artificially raised and lowered by 12%, and then a line was fit to the resulting flux-gradient plot. For simplicity (due to the large number of cases), a linear fit was used in every case rather than trying different fits, though the results of Section 5.2 show that this may not always be the best fit. Using a parabolic fit for a few test cases, instead of a linear fit, did not significantly alter the results.

Unlike gyrokinetic codes, TGLF does not output spatially resolved fluctuation data, so the type of synthetic diagnostic that is applied to the output of gyrokinetic codes will not work for TGLF. Instead, TGLF outputs directly a single fluctuation spectrum, which one can integrate over the wavenumbers that are measurable with the CECE diagnostic. To this end, the ρ_s for the plasma in question was calculated and then compared to the diagnostic spot size in order to determine the extent of the $k_\theta \rho_s$ spectrum that is measurable. This portion of the spectrum was then integrated

to get a total fluctuation level. This method doesn't contain the full complexity of the process that is captured in the synthetic diagnostics for gyrokinetic codes, but it is the only method that works given the outputs of TGLF.

After obtaining the total integrated electron temperature fluctuation level from the output of TGLF, there is one final correction that must be applied for the ASDEX Upgrade discharges (but not for the Alcator C-Mod discharges). As was discussed in Section 5.2, CECE measures the perpendicular temperature fluctuation level. TGLF, however, only outputs the total temperature fluctuation level, as the code cannot separately resolve parallel and perpendicular temperature fluctuations. While the high collisionality on Alcator C-Mod discharges means that these two values are essentially the same, the lower collisionality on ASDEX Upgrade leads to a difference between the two values, by up to 33% [107, 139].

The most rigorous way to account for this difference would be to run nonlinear GENE simulations for every discharge and compute the difference between total and perpendicular fluctuations. This is, however, far too computationally expensive, and would defeat the entire purpose of this new type of multi-discharge validation study. For this reason, the results of Reference [107], presented in Section 5.2, are used as an approximate correction to the output of TGLF, multiplying the output total fluctuation level by 1.33 in order to roughly give the perpendicular temperature fluctuation level. This is clearly an approximation, as the exact difference would depend on the discharge in question, but the discharges considered in this study are similar enough that this method should be at least approximately correct. While the uncertainty on this calculation may be difficult to quantify, as an estimation an additional 5% uncertainty was added to the experimentally calculated fluctuation level uncertainty for the ASDEX Upgrade cases. Despite the caveat, this correction is thought to be more accurate than using the uncorrected level, corroborated by the fact that in test cases, multi-scale TGLF was able to match the corrected fluctuation level, but not

the uncorrected level.

For each experimental case, therefore, many different combinations of TGLF scale configuration, inputs to vary, and applied constraints were used. This resulted in 16 different validation runs for each experimental case with all four constraints available, and 8 different validation runs for experimental cases with three constraints. In total, 168 validation runs were performed.

This validation methodology is now put to use on the 11 experimental discharges described above.

6.3 Results of Multi-Machine Study

This section presents the results of the multi-machine validation study and the resulting criteria for the importance of multi-scale effects. First, the results for the 168 validation runs described in the last section are presented, along with summaries of overall trends and general conclusions that can be drawn from these results. A subset of these results is then used to develop two criteria with which one can determine whether or not multi-scale effects are important for a given plasma. The physical implications of these criteria are also discussed. Finally, TGLF is used to predict temperature profiles for two plasmas on ASDEX Upgrade, one of which satisfies the derived criteria and one of which does not, as an example of the practical implications of these validation efforts.

6.3.1 Validation Results

Table 6.3 summarizes all of the validation results for each of the 168 cases described in the previous section. In this table, a green checkmark indicates that TGLF agreed simultaneously with all of the validation constraints within uncertainty for that particular case, after the inputs for that case had been varied within experimental uncertainty. A red ‘X’ indicates that TGLF failed to match at least one of the constraints

Machine	Discharge	I_p (MA)	P_{aux} (MW)	ρ_{tor}	n_e (10^{19} m^{-3})	Ion-scale						Multi-scale												
						$a/L_{Te}, a/L_{Ti}$			$a/L_{Te}, a/L_{Ti}, a/L_{ne}, Z_{eff}$			$a/L_{Te}, a/L_{Ti}$			$a/L_{Te}, a/L_{Ti}, a/L_{ne}, Z_{eff}$									
Inputs						Q_e	Q_e	Q_e	Q_e	Q_e	Q_e	Q_e	Q_e	Q_e	Q_e	Q_e	Q_e	Q_e	Q_e	Q_e	Q_e	Q_e		
Constraints						Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	Q_i	
						\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	\tilde{T}/T	
						χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	χ_e^{pert}	
Alcator C-Mod	1120706008	0.8	1.2	0.75	6.52	✓	X	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
	1120706017	0.8	1.2	0.75	8.83	✓	X	✓	X	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
	1120706018	0.8	1.2	0.75	10.0	✓	X	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
	1120706019	0.8	1.2	0.75	6.28	✓	✓	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
	1120706030	0.8	4.5	0.75	6.74	X	X	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
ASDEX Upgrade	33585	1.0	0.7	0.75	1.41	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	34301	0.8	0.3	0.39	2.46	X	X	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
	34303	0.6	0.3	0.70	1.58	✓	X	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
	34309	1.0	0.6	0.70	1.85	✓	X	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
	34508	0.6	0.5	0.50	2.47	✓	X	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
	34623	0.6	0.6	0.70	1.80	X	X	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
				0.30	2.91	X	X	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
				0.65	2.50	X	X	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓
				0.27	3.90	X	X	✓	✓	✓	X	✓	X	X	X	X	X	X	X	X	X	✓	✓	✓

Table 6.3: Full validation results for all discharges and conditions considered in this study. A green checkmark indicates agreement between the simulation and experiment within experimental uncertainty. A red ‘X’ indicates disagreement. An empty grey box indicates that the particular set of experimental measurements for that column was not available at that radial location for that discharge. Modified from Reference [65].

within uncertainty. A grey box indicates that the set of validation constraints for a column were not available for that particular plasma condition at that radius. While this binary classification is not particularly precise, it is useful when one is presented with so much data. The agreement with constraints will be quantified later in this section with a validation metric.

This table contains quite a lot of information from which several conclusions can be drawn. First, and most simply, note that when all inputs are varied within uncertainty the multi-scale TGLF model agrees with all available validation constraints within uncertainty for all of the plasma conditions (the furthest right colored cell of each row of Table 6.3). This result is already quite significant, as it successfully validates TGLF in 17 new plasma conditions on two machines, finding good agreement in every case. This is not to say that multi-scale TGLF will always work in every case, as the plasmas considered in this study only cover some finite parameter range. In particular, all of the plasmas are low to moderate power L-mode discharges without internal transport barriers, reversed shear, etc. More complete validation of TGLF will also have to include H-mode, I-mode, high β , and other types of plasmas, as well as looking at other machines besides Alcator C-Mod and ASDEX Upgrade.

In contrast, when one looks at ion-scale TGLF, only some cases where all inputs are varied agree with all available validation constraints (the furthest right colored cell of each row on the ion-scale section of Table 6.3). This indicates that ion-scale TGLF agrees with validation constraints in some cases, but not in all cases, which is consistent with past gyrokinetic validation work showing that multi-scale effects are important in some plasmas but not in all. The task of developing criteria which distinguish plasmas in which multi-scale effects are important from plasmas in which ion-scale simulations are sufficient is exactly the task of determining what differentiates the cases in this table.

More generally, Table 6.3 shows that the multi-scale configuration performs better

than, or equally as well as, the ion-scale configuration in nearly every case. There are a few cases, however, when only temperature gradients are scanned, for which the ion-scale model agrees and the multi-scale model disagrees. In all of these cases the multi-scale model disagrees only slightly outside of uncertainty. These cases are believed to be examples of fortuitous agreement, in which the application of additional validation constraints would eventually show the multi-scale model to be equal or superior to the ion-scale model. These cases reveal the importance of varying more than just temperature gradients, as the multi-scale model always performs better for the cases for which all four inputs are varied.

The effect of adding more validation constraints to the same model and set of input variations is now considered. Additional constraints can by definition only reduce the binary agreement shown in this table, as is seen in a variety of cases for the ion-scale configuration and the reduced input variation multi-scale configuration. These cases for which additional constraints disagree with the experiment, while heat fluxes are found to agree, are examples of what has previously been termed fortuitous agreement, or agreement with a limited set of constraints without truly capturing the plasma turbulent state. This illustrates the importance of including as many validation constraints as possible, and motivates future studies with more than just the four constraints considered here.

Finally, the differences between varying only temperature gradients or varying all four inputs in the validation effort is analyzed. Unsurprisingly, varying more inputs always leads to better or equal agreement with constraints. In particular, some plasmas agree with all available constraints varying only temperature gradients, while for others adding the additional input variations leads to better agreement. This result supports the hypothesis motivated by the work in Section 6.1, that in some cases one must vary more than just temperature gradients to find agreement between simulations and experimental constraints. While such variations would be difficult for

more computationally expensive gyrokinetic simulations, it is certainly worth keeping in mind that this could play an important role in validation results.

A subset of these validation results (those with all inputs varied and all available constraints applied) will now be used to develop criteria for the importance of multi-scale effects, which is one of the main motivations for this study.

6.3.2 Validation Metric and Criteria for the Importance of Multi-Scale Effects

For the purposes of developing criteria for the importance of multi-scale effects, it is necessary to further quantify the difference between the performance of the ion-scale and multi-scale configurations of TGLF. Specifically, this work makes use of a modified version of the Ricci validation metric, χ_{Ricci} [149], which was first introduced in Chapter 5 and is further explained in Appendix D. Recall that χ_{Ricci} is essentially a weighted sum of the differences between the experimentally measured and simulation validation constraint values, normalized by the uncertainty of the experimental measurements.

To minimize the influence of fortuitous agreement, and for the sake of simplicity, only the validation results with all four inputs varied and all available constraints applied to each experimental plasma are used (the furthest right colored cell of each row in the ion- and multi-scale sections of Table 6.3). This section therefore takes the difference between the χ_{Ricci} values for the ion- and multi-scale configurations, with all inputs varied and all available constraints applied, in order to quantify the importance of multi-scale effects:

$$\Delta\chi_{Ricci} = \chi_{ion} - \chi_{multi} \tag{6.1}$$

where χ_{ion} is the Ricci metric calculated for the ion-scale result with all available

constraints and χ_{multi} is the Ricci metric for the multi-scale result with all available constraints.

Using this definition, a value of $\Delta\chi_{Ricci} = 0$ indicates that the ion- and multi-scale models are able to match the validation constraints equally well (though they may do so using different input parameters), and that multi-scale effects are therefore unimportant. A value of $\Delta\chi_{Ricci} = 1.0$ indicates that the multi-scale model performs significantly better than the ion-scale model, and that multi-scale effects are therefore very important. While it is also possible to have negative values of $\Delta\chi_{Ricci}$, indicating that the ion-scale model matches the experiment more closely than the multi-scale model, there is no strong reason for this to be the case outside of small variations in optimization.

There is some discretion in choosing a cutoff value of $\Delta\chi_{Ricci}$, beyond which one says that multi-scale effects are important enough that ion-scale models are no longer sufficient. For the purposes of the work presented in this thesis, a value of $\Delta\chi_{Ricci} = 0.1$ is chosen. Any cases with $\Delta\chi_{Ricci} > 0.1$ are therefore considered to have important multi-scale effects.

Using this quantification of the importance of multi-scale effects as a tool, one can now develop criteria to determine whether or not multi-scale effects will be important without running nonlinear gyrokinetic simulations. Some discussion of possible criteria was already given at the end of Chapter 5, specifically concerning the ratio of linear growth rates at ion and electron scales. This proposal, from References [14, 17, 64], suggested that $\gamma_{high-k}/\gamma_{low-k}$ determined the importance of multi-scale effects, with a cutoff of somewhere between 20 and 40. Recall that γ_{high-k} is the peak linear growth rate in the electron scales ($k_y\rho_s \gtrsim 2.0$), and γ_{low-k} is the peak linear growth rate in the ion scales ($k_y\rho_s \lesssim 2.0$). A limited number of test cases and variation in exactly how each case was treated, however, made a comprehensive analysis difficult for this past work.

Another past study suggested instead that $(\text{MAX}(\gamma/k)_{high})/(\text{MAX}(\gamma/k)_{low})$ should be used as a criterion [22], for which one takes the peak of the linear growth rate spectrum normalized by the wavenumber above $k_y\rho_s > 1.0$, divided by the peak at low wavenumber. Reference [22] suggested that this criterion should have a cutoff value of 1.0. Again, that work was also based on a limited set of discharges.

This thesis tested many possible criteria and selected those which best separated cases with different values of $\Delta\chi_{Ricci}$. The final conclusion is that one must apply two criteria, and that the plasma must satisfy both simultaneously in order for ion-scale simulations to be sufficient. These two criteria, along with their physical interpretations are now presented.

The first of the two criteria is similar to that proposed in Reference [22], though it differs slightly in definition. This thesis defines:

$$\Gamma_{high/low} = (\gamma_{high-k}/k_{high})/(\gamma_{low-k}/k_{low}) \quad (6.2)$$

where again γ_{high-k} and γ_{low-k} are the peak high and low wavenumber linear growth rates, k_{high} is the poloidal wavenumber corresponding to the peak high-k linear growth rate, and k_{low} is the poloidal wavenumber corresponding to the peak low-k linear growth rate. If there is no clear peak at low wavenumber (in the case of a monotonically increasing linear growth spectrum), the linear growth rate at $k_y\rho_s = 2$ was used as γ_{low-k} .

Figure 6-2 shows the results of plotting $\Delta\chi_{Ricci}$ against $\Gamma_{high/low}$ from the 17 plasma conditions described in the last section. With a few exceptions (which are addressed by the second criterion), cases for which $\Gamma_{high/low} < 1.0$ have $\Delta\chi_{Ricci} < 0.1$, and the ion-scale TGLF model is sufficient. On the other hand, as $\Gamma_{high/low}$ increases above 1.0, $\Delta\chi_{Ricci}$ steadily increases to values of approximately 0.7, indicating that multi-scale effects are becoming increasingly important. The outlier on the far right of the graph, which shows less importance of multi-scale effects than one might expect




Figure 6-2: The significance of multi-scale effects, as represented by $\Delta\chi_{Ricci}$, plotted against $\Gamma_{high/low}$, which is the ratio of high and low wavenumber linear growth rate peaks normalized by the wavenumber. Each red square is one validation case. Points in the blue highlighted region, below $\Delta\chi_{Ricci} = 0.1$, are sufficiently represented by ion-scale simulations, while points above this cutoff require multi-scale effects. Originally from Reference [65].

at that value of $\Gamma_{high/low}$, is likely due to fortuitous agreement, and it is hypothesized that adding additional validation constraints would move this point into the more generally observed trend.

In addition to its practical uses for informing future gyrokinetic simulations, the form of this criterion provides information on which physical phenomena are dominantly responsible for the importance of multi-scale effects. As has been mentioned earlier, there are two major sources of differences between ion- and multi-scale simulations. The first is the inclusion of purely electron-scale phenomena, such as electron temperature gradient (ETG) turbulence. These phenomena are also captured in electron-scale only simulations. In addition, multi-scale simulations include the effects of cross-scale coupling between different turbulent modes, moderated by zonal flows [12, 21, 22]. It is this cross-scale coupling that leads multi-scale simulations to often give different results to simple linear additions of heat fluxes from electron- and




Figure 6-3: The significance of multi-scale effects, as represented by $\Delta\chi_{Ricci}$, plotted against $(\gamma_{high-k}/k_{high}^2)/(\gamma_{low-k}/k_{low}^2)$. Each green square is one validation case. This $1/k^2$ normalization separates cases that agree and those that disagree less cleanly than the $1/k$ normalization.

ion-scale simulations.

While it is generally accepted that larger linear growth rates at electron scales should lead to more significant electron-scale contributions to the overall turbulent state of the plasma (as was seen in the criterion in Chapter 5, $\gamma_{high-k}/\gamma_{low-k}$), the relative importance of electron-scale phenomena and cross-scale coupling was unclear. The normalization of $1/k$ in the first criterion, however, indicates that zonal flow mixing and its role in coupling ion- and electron-scales is the most important physical mechanism in determining when multi-scale effects are important [22]. In contrast, a $1/k^2$ normalization would have indicated that the mixing length of a given turbulent mode was dominantly important [21]. This $1/k^2$ normalization is shown in Figure 6-3, revealing significantly less clean separation between cases that agree and those that disagree. In other words, the form of the criterion derived from the validation results above reveals that coupling between ion- and electron-scales, and not the inclusion of purely electron scale transport, is most relevant to the importance of multi-scale simulations. This is consistent with past work, discussed at various points throughout

this thesis, that linearly adding ion- and electron-scale simulations is not equivalent to running a true multi-scale simulation.

The second criterion is also related to cross-scale coupling, but instead of dealing with the strength of electron-scale drives, it concerns the susceptibility of ion-scale turbulence to coupling. Specifically, it is based on the relative drives of ITG (ion temperature gradient) and TEM (trapped electron) modes at ion-scales. This second criterion is defined as:

$$\Sigma_{tem/itg} = \exp\left(\frac{-(\frac{\gamma_{TEM}/k_{TEM}}{\gamma_{ITG}/k_{ITG}} - 1.0)^2}{0.09}\right) \quad (6.3)$$

where γ_{TEM} is the peak linear growth rate of the TEM at ion scales (the TEM can occasionally have a second peak at scales between the ion and electron scales, but this is not of interest here), k_{TEM} is the wavenumber corresponding to the TEM peak, γ_{ITG} is the peak linear growth rate of the ITG, and k_{ITG} is the wavenumber corresponding to the ITG peak.

This criterion is a measure of how closely the peak ITG and TEM growth rates, normalized by wavenumber, compete for dominance at ion-scales. A value of $\Sigma_{tem/itg} = 1.0$ indicates that the two modes are very close, while a value of 0 indicates that one mode dominates. To that end, if either mode is entirely stable, then the value will be approximately 0. Although the final form of this criterion is rather complex, it was constructed specifically to vary between 0 and 1 and so that the cutoff of $\Delta\chi_{Ricci} = 0.1$ occurs at $\Sigma_{tem/itg} = 0.5$. The value of 0.09 in the denominator determines the cutoff, and a cutoff of 0.5 produces the most evenly spaced results.

Figure 6-4 shows the value of both criteria for each plasma condition, plotted as a colored circle representing $\Delta\chi_{Ricci}$ on a two-dimensional plot of $\Gamma_{high/low}$ and $\Sigma_{tem/itg}$. Specifically, turquoise points represent cases where ion-scale simulations are sufficient ($\Delta\chi_{Ricci} < 0.1$), and then green, red, and black represent increasing $\Delta\chi_{Ricci}$ and thus increasing importance of multi-scale effects.




Figure 6-4: A visual representation of both criteria for the importance of multi-scale effects in turbulent transport. Each validation case is plotted on axes of $\Gamma_{high/low}$ and $\Sigma_{tem/itg}$. The color scale shows $\Delta\chi_{Ricci}$, with values below 0.1 (indicating that ion-scale simulations are sufficient) shown in turquoise. The orange shaded regions indicate that multi-scale effects are important due to the breaking of either one or both criteria. Discharges are annotated as follows: ‘C’ for Alcator C-Mod or ‘A’ for ASDEX Upgrade, the last two digits of the discharge number, and (for ASDEX Upgrade cases) an ‘o’ for outer radius or an ‘i’ for inner radius. See Table 6.3 for discharge information. Modified from Reference [65].

This figure reveals that all points in the bottom left quadrant, which satisfy both $\Gamma_{high/low} < 1.0$ and $\Sigma_{tem/itg} < 0.5$, have $\Delta\chi_{Ricci} < 0.1$ and therefore are sufficiently modeled by ion-scale simulations. All points outside of this quadrant have significant multi-scale effects. In particular, note that the two outliers in Figure 6-2, which seemed to elude the first criterion, are in the top left quadrant, and are therefore captured by the second criterion. In general, as one moves further from the bottom left, multi-scale effects become more important, with the exception of the green point in the top right. This point (which was also discussed in Figure 6-2) is likely a case of fortuitous agreement, for which adding more validation constraints would show shortcomings in the ion-scale model.

The physical interpretation of the second criterion also concerns the importance of cross-scale coupling in determining the final turbulent state of the plasma. While the first criterion measures the strength of the electron-scale turbulence drive, and thus its ability to couple to the ion-scales, the second criterion measures the susceptibility of the ion-scale turbulence to coupling with the electron scales. In cases for which there is no clearly dominant mode at the ion-scales (ITG and TEM compete for dominance), then cross-scale coupling, even from weaker electron-scale drive, is important to determining which ion-scale mode becomes dominant, and therefore what the final turbulent properties are.

These criteria are now applied to some of the plasmas used for validation of GYRO and GENE in Chapter 5. This is quite an important comparison, as one of the motivations for developing these criteria is to inform when multi-scale gyrokinetic simulations are necessary, and one should therefore also apply them to validation cases of gyrokinetic codes. First, the criteria are calculate for the L- and I-mode study on Alcator C-Mod from Subsection 5.1.1. Since that study involved global simulations, the metrics are checked at multiple radial locations. For the purposes of this check, only the L-mode portion of the discharge is analyzed. At $\rho_{tor} = 0.88$ (the outer edge of the simulation domain), the plasma has $\Gamma_{high/low} = 0.48$ and $\Sigma_{tem/itg} = 0.90$, which violates the second criterion, indicating that multi-scale effects should be important. At $\rho_{tor} = 0.75$, however, both criteria are satisfied, with $\Gamma_{high/low} = 0.20$ and $\Sigma_{tem/itg} = 0.02$. Interestingly, at the inner edge of the simulation domain, $\rho_{tor} = 0.65$, where the global simulations did quite poorly, both criteria are also satisfied, with $\Gamma_{high/low} = 0.15$ and $\Sigma_{tem/itg} = 0.01$.

At least at first glance, the criteria would seem to correctly predict the disagreement with ion-scale simulations at the outer radius, but incorrectly predict agreement at the inner radius. It is important to note, however, that the values for the criteria just quoted are all for the nominal experimental inputs, which is consistent with how

one would eventually use the criteria, checking for the importance of multi-scale effects before running the nonlinear simulations. If, however, one uses inputs for a heat-flux matched case at $\rho_{tor} = 0.65$, one finds that $\Gamma_{high/low} = 0.50$ and $\Sigma_{tem/itg} = 0.76$, which violates the second criterion. This reveals one possible complication with applying the criteria to plasmas in which the linear growth rates are highly sensitive to small changes in the input gradients. In cases such as this, slightly varying inputs may cause one or the other criteria to be broken, and so one may have to check a few variations of inputs for a complete assessment. It is also possible that cases for which variations of inputs lead to broken criteria are in some sense an edge case, where multi-scale effects may still be marginally important. This serves as a possible explanation for why the ion-scale gyrokinetic simulations did poorly at the inner portion of the domain.

The criteria are next applied to the inner radius GENE validation study on ASDEX Upgrade from Section 5.2.2. For this discharge, 33585 (which is also one of the ones used in the multi-machine TGLF work), at the radial location of that study, $\rho_{tor} = 0.49$, one calculates $\Gamma_{high/low} = 0.59$ and $\Sigma_{tem/itg} = 0.001$. This discharge therefore satisfies the criteria developed in this section, and ion-scale simulations should be sufficient to model the plasma. The same is true when one uses the heat-flux matched input gradients. This is indeed consistent with the GENE validation work, in which GENE was able to simultaneously match all of the available validation constraints: the electron heat flux, the ion heat flux, and the perturbative diffusivity.

To summarize, these two criteria indicate that if either electron-scale turbulence is sufficiently driven, or ion-scale turbulence is particularly susceptible to cross-scale coupling contributions, then multi-scale effects are important and ion-scale simulations will not adequately model the plasma. These criteria have also been applied to two of the gyrokinetic validation studies from Chapter 5, and seem to correctly predict whether or not multi-scale effects will be important.

6.3.3 Profile Prediction

Returning now to one of the original motivations behind all validation work, TGLF will be put to use predicting temperature profiles as an illustration of the implications of the criteria just derived. Profile predictions are done iteratively through the TGYRO framework [158], which was also run within OMFIT [157]. TGYRO runs TGLF at multiple radial locations in the plasma, using the experimental profiles to calculate heat fluxes, and then adjusts the gradients to better match the heat fluxes. These new gradients are integrated to obtain new profiles, for which the process is repeated until the profiles lead to heat flux calculations that match the experimental values. In addition to turbulent transport calculations with TGLF, TGYRO also includes the effects of neoclassical transport, as calculated with the NEO code [159]. The final predicted profile should not depend on the initial profile input (the experimental profile), though starting with the experimental profile rather than an arbitrary profile considerably speeds the convergence process. Depending on the settings used in TGYRO, it is possible that the iterative process will not converge, though that was not the case for any of the plasmas analyzed here. For the work presented here, only the ion and electron temperature profiles were allowed to vary, though TGYRO is also able to vary and predict the density profile.

One difference between this iterative process and that used above in the validation effort is exactly which parameters are allowed to change. The validation effort above worked purely locally, and changed only ion and electron temperature gradients, density gradient, and effective charge (which, for example, also changes collisionality). The TGYRO framework changes electron and ion temperature profiles, though this impacts a variety of other parameters, including ion and electron temperature gradients, T_i/T_e , collisionality, and others. It does not, however, change the effective charge or the density gradient. For this reason, these two processes are similar, but

not exactly the same in their input variation. Finally, since TGYRO is intended for profile prediction, only the final predicted profiles are compared between simulation and experiment. In the end, this is also the comparison that counts if one wants to use simulations to design future machines.

TGYRO was only used to predict profiles between $\rho_{tor} = 0.2$ or 0.3 (inside of which sawtooth crashes dominate transport) and $\rho_{tor} = 0.8$ (outside of which edge and pedestal effects become important). The experimental temperature was therefore used for $\rho_{tor} = 0.8$, and predicted from there inward. For full profile predictions, one must also include a pedestal model (such as EPED [160]) and a core sawtooth model, which is generally time dependent.

Figure 6-5 shows the results of profile predictions for two of the ASDEX Upgrade discharges considered in the validation study above, 33585 and 34623. Discharge 33585 had $\Delta\chi_{Ricci} = 0.02$ at $\rho_{tor} = 0.50$ and $\Delta\chi_{Ricci} = 0.0002$ at $\rho_{tor} = 0.75$. In terms of the derived criteria, at $\rho_{tor} = 0.50$, $\Gamma_{high/low} = 0.18$ and $\Sigma_{tem/itg} = 0.26$. Discharge 34623 had $\Delta\chi_{Ricci} = 0.66$ at $\rho_{tor} = 0.27$ and $\Delta\chi_{Ricci} = 0.60$ at $\rho_{tor} = 0.65$. For the derived criteria, at $\rho_{tor} = 0.27$, $\Gamma_{high/low} = 1.7$ and $\Sigma_{tem/itg} = 0.001$.

The fact that the two discharges predicted here are two of the discharges from the validation study in no way lessens the significance of the agreement or disagreement between the predicted and experimental temperature profiles. The profile predictions and the validation were performed entirely separately, and TGLF was in no way ‘adjusted’ to include the results of the validation study. TGLF’s saturation rule does not include any discharges from ASDEX Upgrade, let alone any of the discharges included in this work. The only reason that predictions were performed on discharges from the validation study is that the experimental work and data analysis required to prepare the code inputs is considerable, and such work had already been performed for all of the discharges in the validation study.

As this figure shows, multi-scale TGLF predictions agree with the experimentally




Figure 6-5: Temperature profile predictions for two ASDEX Upgrade discharges. The left column in discharge 33585 and the right is 34623. The top row is electron temperature and the bottom is ion temperature. Experimental profiles are shown in turquoise, with uncertainty represented by the shaded region. Multi-scale TGLF predictions are shown in green and ion-scale in red. Only radii inside of the edge and outside of the sawtooth region are shown. Figure modified from Reference [65].

measured ion and electron temperature profiles within uncertainty for both discharges, with only a slight deviation in the electron temperature of discharge 34623 around $\rho_{tor} \approx 0.7$. It is likely that variation of the density gradient or effective charge would be able to resolve this discrepancy. Ion-scale TGLF, on the other hand, does reasonably well for 33585, with only slight deviations in both the ion and electron temperatures, but severely over-predicts the electron temperature in discharge 34623. These predictions are robust to the spatial resolution of the prediction and to the edge pinning location (moved in or out by $0.05 \rho_{tor}$).

The discrepancy between the ion-scale model and experiment in discharge 34623 and the agreement in discharge 33585 is consistent with the values of the criteria for both discharges. In 33585, both criteria are met, indicating the ion-scale simulations should be sufficient, while in 34623, $\Gamma_{high/low} = 1.7 > 1.0$, indicating that multi-scale effects are important and that one should not trust the result of ion-scale simulations.

Conclusions and Implications of Multi-Machine Validation Study

This chapter has presented a validation effort with slightly different goals than those presented in Chapter 5, focusing on developing more general criteria to determine the importance of multi-scale effects, rather than focusing on validating a specific code for one particular plasma discharge. The results of this work were manifold, starting with the validation of multi-scale TGLF for 17 plasma conditions from Alcator C-Mod and ASDEX Upgrade, with three or four constraints for each condition. This is itself a significant result, as it shows that TGLF is a robust transport model, which agrees with all available experimental measurements over a wide range of L-mode plasmas on two different machines.

The validation portion of the study also further proved the importance of including as many validation constraints as are available, in order to avoid fortuitous agreement, which was seen for several cases when only heat fluxes were compared. In addition, this effort varied many input parameters, beyond just temperature gradients, revealing that this input variation in some cases leads to validation agreement where one would have otherwise found disagreement.

In all, the validation portion of this study represents a new methodology, focusing on breadth of validation, rather than depth. This type of study will be necessary in the future, as it is vital to validate turbulent transport models in as many different plasma conditions as possible before using them predictively. Continuation of the study presented here would involve additional types of plasmas (H-modes, I-modes,

etc.), additional machines (DIII-D, JET, etc.), and additional constraints (density fluctuations, etc.). Validation in only a few plasma conditions or on only a single machine, even very thoroughly, represents insufficient evidence to build confidence in future predictions for future machines.

More generally, this chapter used the results of the validation study to investigate the physics of multi-scale turbulence effects, and to develop criteria that can be used to determine when multi-scale simulations are necessary. The two criteria, $\Gamma_{high/low}$ and $\Sigma_{tem/itg}$, both deal with the importance of cross-scale coupling in determining the final turbulent state of the plasma. The form of the first criterion reveals that it is cross-scale coupling of turbulence, moderated by zonal flows, which is dominantly responsible for the importance of multi-scale effects. Sufficiently high electron-scale linear growth rates will lead to coupling with the ion-scales. The second criterion shows that even if electron-scale drives are weak, if the ion-scale turbulence has no clearly dominant mode, then cross-scale coupling is still important, and multi-scale simulations are still necessary. These criteria were also applied to two of the gyrokinetic studies from Chapter 5, correctly predicting whether or not the ion-scale simulations would agree with validation constraints.

Finally, the practical implications of validation were illustrated by using ion- and multi-scale TGLF to predict temperature profiles for two ASDEX Upgrade discharges, one which satisfied both criteria and one which violated the first criterion. The severe overprediction of the electron temperature profile for discharge 34623 with ion-scale TGLF illustrates exactly why one should not trust the results of models which have not been thoroughly validated. Such a prediction would lead one to be overly optimistic as to the performance of a future device. The profile predictions also show that an ion-scale model does quite well for a plasma which satisfies both of the criteria developed in this chapter.

Chapter 7

Conclusions

This thesis has described a wide range of work, from new experimental hardware to new analysis techniques and simulation results, but the central narrative throughout has been validation of turbulent transport simulations with the hope that they will eventually be ready to predict the performance of future fusion machines. This final chapter summarizes the work presented in this thesis, as well as suggesting future research endeavors motivated by the results of this work.

The Work of this Thesis

The first two major contributions of this thesis were the development of analysis techniques and hardware in order to experimentally measure two validation constraints for turbulent transport models, the electron perturbative thermal diffusivity and electron temperature fluctuations. Chapter 3 described in detail the method of measuring the perturbative thermal diffusivity in tokamak plasmas by tracking the propagation of heat pulses generated by partial sawtooth crashes. This new method avoids the non-diffusive heat transport associated with full sawteeth and does not require active perturbation of the plasma with modulated electron cyclotron heating, requiring only a plasma that naturally produces partial sawteeth. It was also checked against mea-

surements from the more established modulated electron cyclotron heating method, finding good agreement. The perturbative diffusivity measured in this fashion is an essentially free additional validation constraint in many naturally sawtoothed plasmas, and for this reason nearly all of the data in that chapter was from previously run plasmas, requiring no additional experimental run time. The requirement for sawtooth crashes may exclude some particularly high performance plasmas with a q greater than 1 on the magnetic axis. As was shown in Chapters 5 and 6, the perturbative diffusivity is a valuable validation constraint when assessing the physical accuracy of turbulent transport models. In addition, perturbative diffusivity measurements from more than 70 discharges on both Alcator C-Mod and ASDEX Upgrade were used to investigate the governing physics of heat pulse propagation, revealing a strong correlation with collisionality across both machines.

Chapter 4 described the development of both hardware and analysis techniques for the correlation electron cyclotron emission diagnostic, used to measure electron temperature fluctuations. This thesis work constructed intermediate frequency electronics sections for CECE diagnostics on both Alcator C-Mod and ASDEX Upgrade. In particular, the hardware on ASDEX Upgrade represented a significant improvement over previous systems, and has taken arguably some of the cleanest CECE data on any tokamak to date. New analysis techniques were also developed analytically in order to more accurately calculate temperature fluctuation levels from the raw CECE radiometer data. These techniques were thoroughly tested with both synthetic and experimental data. In terms of purely experimental observations, these CECE diagnostics were used to investigate I-mode physics on Alcator C-Mod and H-mode physics on ASDEX Upgrade. They also provided key validation data to the studies in Chapters 5 and 6.

Having developed two experimental tools necessary for thorough validation of turbulent transport models, this thesis then moved on to two chapters describing

directed validation studies. Chapter 5 described validation of two gyrokinetic codes, GYRO and GENE. The validation studies in this chapter followed a fairly established methodology, looking at a single code in one or two plasmas in great depth. The first study compared GYRO simulations of L- and I-mode plasmas on Alcator C-Mod (using as validation constraints the electron heat flux, ion heat flux, and both measurements described in this thesis), finding that the code agreed rather poorly with experimental measurements from both plasmas. One possible reason for this disagreement was the absence of multi-scale physics in the GYRO simulations, a hypothesis motivated by the second validation study of the chapter, which showed that in at least one case, multi-scale simulations resolved discrepancies between ion-scale simulations and experimental observations. The next two validation studies in the chapter focused on validating GENE on ASDEX Upgrade. Both studies, which applied different validation constraints to two radii in the same plasma discharge, found reasonably good agreement between ion-scale simulations and experiment. This raised the question of when ion-scale simulations were sufficient to model the plasma turbulence, and when multi-scale simulations were necessary.

The next chapter in the thesis, Chapter 6 attempted to answer this question by using the quasi-linear gyrofluid code, TGLF. First, multi-scale TGLF was applied to the same L- and I-mode plasmas analyzed with GYRO in Chapter 5. Contrary to the previous hypothesis, however, multi-scale TGLF also disagreed with experimental measurements, suggesting that some other effect may be responsible for the disagreement. One possibility was that one may have to vary more than just the input temperature gradients in order to find a match with experimental measurements. All of these results motivated the design of the final validation study in the thesis. Instead of focusing on a single plasma on one machine, one code, TGLF, was applied equally to 11 discharges from 2 machines, all within the parameter range for which TGLF's saturation rule was tuned. TGLF was run varying up to four inputs through the VI-

TALS framework, allowing investigation of the importance of varying multiple inputs. TGLF was also run in both ion- and multi-scale configurations in order to investigate when multi-scale effects were important and when they were not. This final study involved 168 combinations of discharge, location, constraints, inputs, and scale setting. The results of the study were twofold. First, multi-scale TGLF was shown to agree with all available validation constraints for all 17 plasma conditions analyzed, once all four inputs were allowed to vary within uncertainty. Second, two physical criteria to predict whether or not multi-scale interactions would be important were developed, offering insight into the physical phenomena behind the coupling of electron- and ion-scales. This final validation effort brought together all that was learned throughout the entire thesis to produce both practical and fundamental physics results.

It is worthwhile to briefly consider some larger picture implications of these validation results. First, as was mentioned before, there is really no reason to run ion-scale TGLF outside of targeted physics investigations (such as the one performed in this thesis), as multi-scale TGLF more closely matches experiment in all cases and requires negligibly more computation time. More generally, this work indicates that for cases which violate either criterion from Chapter 6, one cannot in good faith use results from any ion-scale gyrofluid or gyrokinetic simulation, as multi-scale interactions are clearly important. This is perhaps a controversial statement, and comes with the caveat that ion-scale simulations are still useful for answering certain specific physics questions. These validation results strongly indicate, however, that for cases which violate either criterion, ion-scale simulations are not an accurate representation of reality. This conclusion by itself is incredibly strong motivation for increased allocation of computing resources to multi-scale gyrokinetic simulations, and then the use of these simulations to develop quasi-linear saturation rules that span larger parameter spaces.

In summary, this thesis has developed analysis techniques and hardware for two

validation constraints and then applied these constraints, and others, to validation of three turbulent transport codes on two tokamaks. The hardware and analysis techniques have been documented in several publications and will hopefully be used in future validation studies by other researchers. In addition to the specific validation results presented in this thesis, which add to the continually growing collective knowledge of the validation community, this thesis answered a concrete question concerning the importance of multi-scale effects, resulting in both practical criteria and physics understanding. The final work with TGLF also pioneered a new validation study methodology, which will hopefully increase the breadth of validation work in the future.

The Road Ahead

While this thesis has contributed substantially to the ongoing work to validate turbulent transport models, and more broadly to the grand march toward fusion energy, there is of course, much remaining to be accomplished. Fortunately, the work in this thesis points to several concrete next steps.

Regarding the partial sawtooth perturbative diffusivity measurements, it would be greatly beneficial to apply this measurement to additional machines beyond Alcator C-Mod and ASDEX Upgrade. This would allow continued expansion of the dataset used to analyze the collisionality dependence, and would set precedent for then using this measurement as a validation constraint on those machines as well. In addition, while this thesis did perform a comparison with modulated electron cyclotron heating measurements, it would be worthwhile to perform additional comparisons on other discharges to further confirm that the two methods give the same result. Finally, all of the analysis in this thesis relied on manually identified partial sawtooth crashes. Development of a code to automatically identify partial sawtooth crashes and then perform the perturbative diffusivity analysis would go a long way toward making this

validation constraint more mainstream.

The CECE results in this thesis also suggest a number of hardware improvements and open questions to be addressed. First, it would certainly be beneficial to install CECE diagnostics on additional machines, beyond the two machines in this thesis and the few other systems in the world (TCV and DIII-D). Additional machines would allow measurements in an expanded parameter regime, and would also contribute to making electron temperature fluctuation measurements a standard validation constraint. In constructing a new system, several hardware lessons learned from this thesis work may be applied. First, the careful attention to reducing noise through separating signal and power grounds and reducing spurious cable length definitely paid off, and should be part of any future CECE, or more broadly, any microwave diagnostic, design. In addition, during operation of the diagnostic, it became clear that with a large number of fixed channels in a comb, the tunable YIG filters were not used very often. For this reason, a new system would likely replace the YIG filters with more fixed channels. The YIGs are both more expensive and take up more space than the fixed channels, not to mention being more complicated due to the required control circuit, so removing one YIG channel could enable the addition of two or three more fixed channels, further expanding the comb coverage. At least as the system on ASDEX Upgrade was operated, this would lead to more useful data. There is of course a trade off in flexibility, and so for a different set of research goals, other systems may wish to retain the YIG channels.

In terms of physics investigations with CECE, two promising experimental research directions were mentioned in Chapter 4. Additional work in both of these areas may deepen understanding of how core temperature fluctuations change in high performance regimes, and of the relationship between core turbulence and increased plasma performance. First, some example measurements were shown from low current I-modes on Alcator C-Mod, which did not show the characteristic reduction in

core temperature fluctuations that is typically observed after the transition from L-mode to I-mode. Further investigation of these low current I-modes may help shed light on what it is about the standard I-mode transition that leads to reduced fluctuation levels in the plasma core and improved confinement. Similarly, data was shown from low current H-modes on ASDEX Upgrade, with and without resonant magnetic perturbations, revealing a change in outer core temperature fluctuations even though RMPs are thought to mainly affect the plasma edge. Investigating the effect of RMPs on electron temperature fluctuations may also shed light on their role in stabilizing edge localized modes and how one might develop other stationary high performance regimes. Finally, while CECE has mostly been used so far to investigate core turbulence, it is also possible to observe closer to the plasma edge, though one must be careful about optical thickness. In particular, there has recently been a lot of interest in the weakly coherent mode near the edge of I-mode plasmas. By combining CECE measurements with density fluctuation measurements, one may gain insight into the decoupling of heat and particle transport in the I-mode pedestal.

Finally, the validation work in this thesis also suggests a number of future directions. First, it would be valuable to eventually run fully multi-scale (with a realistic electron to ion mass ratio) GENE simulations, and to compare these results to multi-scale GYRO and TGLF runs. This would bolster confidence in the multi-scale gyrokinetic results that have already been published, and ensure that the two codes give similar predictions. Along the same lines, additional multi-scale gyrokinetic runs are required in order to generate better saturation rules for quasi-linear codes like TGLF, and expand the parameter range in which they are valid.

Specific to the results in this thesis, it would be good to use multi-scale gyrokinetic simulations to check a few of the TGLF cases presented in the last study to ensure their accuracy, as well as to check the application of the derived criteria for more than just a few gyrokinetic cases (these comparisons were not performed as part

of this thesis due to the significant computational resources required for multi-scale gyrokinetic simulations). The last validation study, while already much broader than most past work, should eventually be expanded to include even more machines and plasma discharges, specifically higher performance regimes such as I- and H-mode. If at some point the study finds parameter space in which multi-scale TGLF is unable to match constraints, it may lead to substantive improvements to the code. For example, one could run multi-scale gyrokinetic simulations with these parameters, and then include those runs in a new saturation rule. A broad study such as this may also reveal other important physics that should be included in turbulent transport simulations.

This last study would also benefit from including additional validation constraints, beyond the four used in this thesis. It could also expand to considering both particle and momentum transport, in addition to the heat transport that was the focus of this thesis. Such a study could include density fluctuations as a validation constraint, made with, for example, reflectometry (described in Appendix E) and high time resolution interferometry (described in Appendix G). Finally, one could apply more rigorous statistical analysis to a larger set of validation data from this type of study, including perhaps various Bayesian statistical methods, in order to quantify the improvement in model agreement as the number of free parameters (varied inputs) is increased.

More generally, turbulent transport models represent just one research focus in the broader fusion research field. This research must continually interact with other research topics, such as the study of plasma heating, edge and pedestal physics, and divertor physics, in order to develop a physical understanding and predictive capability for the entire tokamak plasma. By combining knowledge from all of these regions of the plasma, and by developing integrated models of the entire machine (an example of which is described in Appendix F), one will hopefully be able to

inform engineering decisions when designing future machines, and also indicate which engineering research paths are most important to achieve the goal of fusion energy.

The path to the clean, sustainable, and safe energy source promised by magnetic fusion has been long, and fraught with seemingly insurmountable obstacles. This thesis has made its contribution to the march toward that goal, but significant challenges remain. In spite of all of this, researchers around the world dedicate themselves to making fusion energy a reality. It is too important not to try. And that is why we do.

Appendix A

History of Perturbative Thermal Diffusivity Measurements

The work in this appendix was originally presented as the Appendix of Reference [55].

The perturbative diffusivity discussed in Chapter 2 and 3 has a fairly long history of measurement, which may be relevant to the reader interested in how the method presented in this thesis fits into the larger historical context.

Using heat pulses to deduce the diffusivity in tokamak plasmas has a long history, and due to its natural presence in many tokamak plasmas, the sawtooth crash was one of the first sources used for this purpose. Tracking heat pulses generated by sawtooth crashes was first performed on the ORMAK tokamak in 1977 [57]. This early work calculated the thermal diffusivity via the time-to-peak method, which tracks the propagation of the peak of the heat pulse [57]. For comparison, the Extended-Time-to-Peak method used in this thesis also includes information from the amplitude of the heat pulse as it moves radially, leading to considerably more accurate calculations. It wasn't until roughly 1986 that modulated electron cyclotron heating (ECH) was used to generate heat pulses [161]. Some work has also tracked 'cold' pulses generated by edge impurity injection [162], though this method is not a focus of this thesis.

These first measurements found that the diffusivity calculated by tracking sawtooth generated heat pulses exceeded the power balance diffusivity by factors of up to 15 [57], but did not at the time recognize that the perturbative and power balance diffusivities were measuring different aspects of the plasma. Some later work on ORMAK partially resolved the discrepancy by using more detailed calculations, but a discrepancy remained [163].

Similar methods were then applied to other tokamaks, with varied results. Work on the ISX-B tokamak found agreement between the power balance and ‘heat pulse’ diffusivities [164], while work on Doublet III [165] and TFTR [166] found that the heat pulse measured diffusivity significantly exceeded the power balance diffusivity. Around this time a method of calculating the heat pulse diffusivity based on Fourier analysis of sawtooth heat pulses was also developed, and was applied to Doublet III [165] and TFTR [166], but did not show significantly different results to the earlier calculation methods. The observations from these larger machines led some to believe that the power balance diffusivity was transiently enhanced during the propagation of heat pulses [166].

Further investigation of the discrepancies between the heat pulse diffusivity and the power balance diffusivity led to the theory that they were in fact different physical properties, one governing the propagation of perturbations and one governing the steady state equilibrium of the plasma [54]. This is reflected in the different definitions given in Chapter 2, and the graphical illustrations in Figure 2-2.

This realization was also accompanied by the development of the Extended-Time-to-Peak calculation method used in this thesis [84]. Later work also compared the Extended-Time-to-Peak and Fourier calculation methods, finding agreement within uncertainty for several experimental conditions [60, 85]. Using these methods, several studies on JET found agreement between the power balance diffusivity and the perturbative diffusivity measured with sawtooth heat pulses [54, 61].

Still, the realization that the power balance and perturbative diffusivities were different physical quantities did not entirely dissipate concerns over the large discrepancy between power balance and perturbative diffusivity measured with full sawtooth heat pulses, by up to a factor of 20 on TFTR [54]. Subsequent investigation revealed that the sawtooth crash itself was a possible cause for this discrepancy, since the sawtooth event led to additional enhanced transport even beyond the mixing radius, likely caused by the macroscopic MHD perturbations to the plasma equilibrium, which was termed “ballistic” transport [167]. This theory was supported by calculations suggesting that the ballistic transport was non-diffusive, and so would lead to large enhancements if one tried to calculate an effective diffusivity [167]. This work therefore suggested that full sawtooth crash-generated heat pulses were unsuitable for calculating the perturbative diffusivity, as the sawtooth event led to non-diffusive transport even far outside of the mixing radius.

This result did not immediately end perturbative diffusivity calculations based on full sawtooth crashes, and for the next decade or so, various studies debated the validity of using full sawtooth heat pulses, including the possibility that one could use the measurement as long as certain spatial or temporal limits were respected [62, 168, 169, 170]. Reference [53] gives a thorough review of perturbative diffusivity measurements based on heat pulse propagation up to this point.

Further evidence emerged from both TFTR and DIII-D in 2000, however, which mostly ended the use of full sawtooth crash-generated heat pulses to measure perturbative diffusivity [58]. This work contained detailed comparisons between experimental data and computational models, revealing that the shape and propagation of sawtooth crash heat pulses were inconsistent with a diffusive process, even if one included a strong nonlinear dependence of the diffusivity on the temperature gradient. This same work was the first to suggest that partial sawteeth might pose a viable alternative, showing that partial sawtooth heat pulses were indeed consistent with a

diffusive model. Finally, it was shown that the ballistic transport extended to at least twice the inversion radius in some cases, leaving very little radial range over which the propagation of full sawtooth heat pulses might have been considered diffusive. The results of this work were later confirmed and expanded upon by other studies [171].

Since this work, it has generally been accepted that ballistic transport makes the use of full sawtooth heat pulses unreliable for measuring the perturbative diffusivity [56, 172], though there are a few examples of studies that have continued to use that method [173, 174]. It is for this reason that before the publication of the work performed in this thesis, modulated ECH had been the preferred source of heat pulses for measuring the perturbative diffusivity [56].

Turning now to the physics behind these conclusions, the reason that the ballistic transport is thought to be non-diffusive is that it is heavily influenced by the macroscopic MHD physics of the sawtooth crash [58]. References [83, 175, 176] discuss the physics of the sawtooth crash in considerable detail, including the fast redistribution of energy and particles which preclude the use of full sawteeth for measuring a diffusivity.

While the partial sawteeth described in this thesis (which are also sometimes called ‘compound’ sawteeth in the literature), are caused by the same fundamental physical processes as full sawteeth, their considerably smaller radial extent and amplitude mean that they are not accompanied by the same ballistic transport as full sawteeth [58, 83]. Qualitatively, one can also show that partial sawteeth result in a considerably smaller redistribution of energy than full sawteeth, using techniques such as those in Reference [177]. The interested reader could also apply more quantitative techniques, such as those in Reference [178].

Appendix B

Calculation of Numerical Prefactor for Perturbative Thermal Diffusivity

Chapter 3 discussed in detail the use of partial sawtooth crash-generated heat pulses in order to measure the perturbative diffusivity in the plasma. In particular, this thesis used the Extended-Time-to-Peak calculation method, shown in Equation 3.1 [54, 84]. For convenience, this is repeated here:

$$\chi_e^{pert} = 4.2a_c \frac{v_{HP}}{\alpha} \quad (\text{B.1})$$

where:

$$v_{HP} = \sqrt{\kappa} \frac{a}{a-s} \left(\frac{dt_{peak}}{dr} \right)^{-1} \quad (\text{B.2})$$

is the radial velocity of the peak of the heat pulse (in m/s) and:

$$\alpha = 10(a-s) \frac{d}{dr} \log(A) \quad (\text{B.3})$$

is a parameter describing damping of the heat pulse as it propagates radially (unitless).

The aim of this appendix is to derive Equation 3.1, including the numerical prefactor of 4.2. This derivation was originally performed in References [54, 84], but in particular this thesis wanted to confirm that the numerical prefactor of 4.2 would not change depending on the machine to which the method was applied. This appendix therefore starts by repeating the analytic derivation of Equation Equation 3.1 from References [54, 84], and then describes the numerical calculations performed to get the prefactor of 4.2.

Following Reference [54], this derivation starts with a simplified version of Equation 2.6, the fluid equation describing energy flux in the plasma, for electrons. Specifically, consider only:

$$\frac{3}{2} \frac{d}{dt}(n_e T_e) \approx -\nabla \cdot \vec{q}_e + S_{E,e} \quad (\text{B.4})$$

where, as before, n_e is the electron density, T_e is the electron temperature, \vec{q}_e is the electron heat flux, and $S_{E,e}$ is the energy source term (which in this derivation also includes the contribution of electron-ion collisions, C_s). Since this derivation is focused only on the dominant perturbed part of the equation, convection, compression, and steady plasma heating will be ignored.

Now, split each variable into equilibrium and perturbed parts, $x = x_{eq} + \tilde{x}$, and take only the perturbed part of the equation. Assuming only a temperature heat pulse (the density is unperturbed), one gets:

$$\frac{3}{2} \frac{d}{dt}(n_e \tilde{T}_e) = -\nabla \cdot \vec{\tilde{q}}_e + \tilde{S}_{E,e} \quad (\text{B.5})$$

If one then applies a diffusive model to the heat transport in this situation, one can write:

$$\vec{q}_e = -n_e \chi_e \nabla T_e \quad (\text{B.6})$$

If one then allows χ_e to be a function of T_e and ∇T_e , focuses on transport only in the radial direction ($\nabla \rightarrow \partial/\partial r$), and limits oneself to first order in perturbed quantities, one can write:

$$-\tilde{q} = n_e \tilde{\chi}_e \nabla T_e + n_e \chi_e^{PB} \nabla \tilde{T}_e \quad (\text{B.7})$$

where χ_e^{PB} is the same quantity defined in Chapter 2, and is the steady state thermal diffusivity.

Now, define:

$$\chi_e^{pert} = -\frac{1}{n_e} \frac{\partial Q_e}{\partial \nabla T_e} = \chi_e^{PB} + \frac{\partial \chi_e}{\partial \nabla T_e} \nabla T_e \quad (\text{B.8})$$

as in Equation 2.22.

Plugging everything into Equation B.7 and staying at first order in the perturbation gives:

$$-\tilde{q} = n_e \chi_e^{pert} \nabla \tilde{T}_e + n_e \left(\frac{\partial \chi_e}{\partial T_e} \right) (\nabla T_e) \tilde{T}_e \quad (\text{B.9})$$

One can show that for any reasonable explicit temperature dependence in χ_e , the second term will have a negligible effect on the calculation of the perturbative diffusivity and is therefore dropped from the remainder of this derivation [54].

Plugging this into Equation B.5 gives:

$$\frac{3}{2} n_e \frac{\partial \tilde{T}_e}{\partial t} = \nabla \cdot (n_e \chi_e^{pert} \nabla \tilde{T}_e) - \frac{n_e}{\tau_d} \tilde{T}_e \quad (\text{B.10})$$

where $\tilde{S}_{E,e}$ has been assumed to be linear in \tilde{T}_e to first order, with a characteristic damping time constant:

$$\tau_d = -n_e \left(\frac{\partial S_{E,e}}{\partial T_e} \right)^{-1} \quad (\text{B.11})$$

This assumes that radiative losses are much smaller than collisional losses, which is the case for nearly all high performance tokamak plasmas. Note that the standard time-to-peak analysis implicitly assumes that $\tau_d \rightarrow \infty$. Allowing this to be finite is what differentiates the extended-time-to-peak method.

Following now Reference [84], it is assumed that τ_d is a constant, uniform in space and time over the propagation region of the heat pulse. This is not true in a real tokamak plasma, as the collisional portion of the damping will certainly depend on the background profiles, but the approximation allows the diffusive calculation carried out here.

The relative strength of the damping term in Equation B.10 is characterized with the dimensionless quantity:

$$K = \frac{a^2}{\tau_d \chi_e} \tag{B.12}$$

where a is the machine minor radius.

One now proceeds to solve Equation B.10 numerically, by imposing an initial temperature perturbation onto background profiles in a cylindrical geometry, and allowing the perturbation to propagate radially. By calculating v_{HP} and α , based on their definitions in Equations B.2 and B.3 respectively, and then comparing them to the input value of χ_e^{pert} , one is able to derive the relationship in Equation 3.1, including the prefactor. These calculations are now described in greater detail.

The machine size and perturbative diffusivity as inputs. In addition, one must choose the shape and location of the initial perturbation, the effect of which is discussed in greater detail below.

When one runs these simulations, one finds that the dimensionless parameters av_{HP}/χ_e^{pert} and α both depend linearly on K , as seen in Reference [84] and in the

repeated calculations in this thesis. They are thus parameterized as follows:

$$\frac{av_{HP}}{\chi_e^{pert}} = p_v + q_v K \quad (\text{B.13})$$

and

$$\alpha = p_\alpha + q_\alpha K \quad (\text{B.14})$$

Solving for χ_e^{pert} gives:

$$\chi_e^{pert} = c_1 \frac{av_{HP}}{\alpha + c_2} \quad (\text{B.15})$$

where $c_1 = q_\alpha/q_v$ and $c_2 = c_1 p_v - p_\alpha$.

Finally, assuming $c_2 \ll \alpha$, which is the case in all of the calculations performed in Reference [84] and in this thesis, one can write:

$$\chi_e^{pert} \approx \gamma a_c \frac{v_{HP}}{\alpha} \quad (\text{B.16})$$

where $a_c = a\sqrt{k}$ now includes a correction for an elongated plasma and:

$$\gamma = c_1 \left(1 - \frac{c_2}{\alpha}\right) \quad (\text{B.17})$$

This γ is the prefactor in Equation 3.1. The average α for the set of simulations is used in this relationship.

To summarize the workflow, one sets up background equilibrium profiles and an initial perturbation in a cylindrical geometry. One then solves Equation B.10 for a given χ_e^{pert} numerically for multiple values of the dimensionless parameter K (which sets τ_d). One then calculates v_{HP} and α for each of the input K values. Plotting av_{HP}/χ_e^{pert} and α against K , one performs linear fits to find p_v , q_v , p_α , and q_α . These are used to calculate γ , the numerical prefactor in Equation 3.1. The published

value of this prefactor in Reference [54] is 4.2. In the work here, the calculation was performed in MATLAB, the results of which are briefly mentioned in Reference [55].

The resulting prefactor may depend on the input perturbation location and shape, as well as the location at which one measures the perturbative diffusivity. In order to check the effect of these inputs on the value of the prefactor, a number of scans were run. Using parabolic, gaussian, triangle, and even square initial perturbations had less than a 1% effect on the calculated prefactor, as long as the measurement location was at least $0.1a$ beyond the initial perturbation location. Similarly, using a variety of experimentally relevant perturbation widths resulted in variations of less than 10% in this prefactor, once the pulse had propagated $0.1a$ radially outward.

Finally, as one varied the exact location of the initial perturbation (from $r/a = 0.1-0.5$) and measurement location (from $r/a = 0.3-0.9$), such that the measurement location was always $0.1a$ further out than the location of the initial perturbation, the calculated numerical prefactor did vary between 3.8 and 4.3. This variation of less than 10% is consistent with the published uncertainty of 10%, and is less than the experimental uncertainty in every case analyzed in this thesis, which is usually at least 15%. It would be possible to recalculate this prefactor for the initial perturbation location and measurement location of every experimental measurement, but because the variation is less than the experimental uncertainty, the slight improvement to the accuracy is likely not worth the additional work. For this reason, the published value of 4.2 is used in Equation 3.1 throughout this thesis.

This completes the derivation of Equation 3.1, including its numerical prefactor.

Appendix C

Derivation of CECE Equations

This appendix derives four equations presented in Chapter 4 that are used in analysis and interpretation of CECE data: Equation 4.1, which calculates the total temperature fluctuation level (derived in both the frequency and time domains); Equation 4.8, which calculates the uncertainty on this value; Equation 4.9, which states the sensitivity limit of CECE measurements; and Equation 2.29, which is used to calculate the optical depth of the second harmonic X-mode ECE.

C.1 Frequency Domain Derivation of Total Temperature Fluctuation

First, this appendix will derive Equation 4.1 in the frequency domain, using the definition of γ_c from Equation 4.2, highlighting differences with previous derivations that have led to Equation 4.5. This derivation was originally published in Reference [81].

Consider one signal collected by an ECE radiometer, $x(t)$, which includes contributions from the steady state temperature T_e , temperature fluctuations $\tilde{T}_e(t)$, and

thermal noise $\tilde{\mathcal{N}}_x(t)$ [71, 120].

$$x(t) = c_x(T_e + \tilde{\mathcal{T}}_e(t))(1 + \tilde{\mathcal{N}}_x(t)) \approx c_x(T_e + \tilde{\mathcal{T}}_e(t) + T_e\tilde{\mathcal{N}}_x(t)) \quad (\text{C.1})$$

where c_x is a calibration factor that relates the radiometer signal to the absolute temperature. This factor will cancel out later in this derivation, so absolute calibration is not necessary for CECE operation using this calculation method. It is assumed in this derivation that $\tilde{\mathcal{T}}_e \ll T_e$ and $\tilde{\mathcal{N}}_x \ll 1$, so that the quadratic term $\tilde{\mathcal{T}}_e\tilde{\mathcal{N}}_x$ is negligibly small [71].

One then takes the Fourier transform of Equation C.1, defining $X(f)$ to be the Fourier transform of $x(t)$, $\tilde{\mathcal{T}}_e(f)$ to be the Fourier transform of $\tilde{\mathcal{T}}_e(t)$, and $\tilde{\mathcal{N}}_x(f)$ to be the Fourier transform of $\tilde{\mathcal{N}}_x(t)$. This gives:

$$X(f) = c_x(\delta(f)T_e + \tilde{\mathcal{T}}_e(f) + T_e\tilde{\mathcal{N}}_x(f)) \quad (\text{C.2})$$

where $\delta(f)$ is the Dirac Delta Function, indicating that the background temperature leads to a constant component at zero frequency.

This derivation is concerned with only the fluctuating part of the frequency spectrum ($f \neq 0$) such that one can ignore the equilibrium temperature. Consider now the cross-spectral density function [115] of two signals x and y :

$$G_{xy}(f) = \langle X(f)^*Y(f) \rangle \quad (\text{C.3})$$

where the triangle brackets represent ensemble averaging and the asterisk represents the complex conjugate, and the auto-spectral density function of x :

$$G_{xx}(f) = \langle X(f)^*X(f) \rangle \quad (\text{C.4})$$

Inserting Equation C.2 into Equation C.3 gives:

$$\begin{aligned}
G_{xy}(f) = c_x c_y & \left(\langle \tilde{T}_e^*(f) \tilde{T}_e(f) \rangle + \cancel{T_e \langle \tilde{N}_x^*(f) \tilde{T}_e(f) \rangle} \right. \\
& \left. + \cancel{T_e \langle \tilde{T}_e^*(f) \tilde{N}_y(f) \rangle} + \langle T_e^2 \tilde{N}_x^*(f) \tilde{N}_y(f) \rangle \right)
\end{aligned} \tag{C.5}$$

Since the thermal noise in each channel is uncorrelated with the turbulent temperature fluctuations, the $\langle \tilde{T} \tilde{N} \rangle$ terms are dropped, as in References [71, 73] and others (even if correlated, these terms would be quadratic in small parameters). Note that ensemble averaging brackets will be dropped in the remainder of this derivation to minimize notational clutter. This then reduces to:

$$\begin{aligned}
G_{xy}(f) &= c_x c_y \left(\tilde{T}_e^2(f) + T_e^2 \tilde{N}_{xy}^2 \right) \\
&= c_x c_y \tilde{T}_e^2(f) + G_{noise}
\end{aligned} \tag{C.6}$$

This assumes that the thermal noises \tilde{N}_x and \tilde{N}_y are mostly uncorrelated, except for some small part \tilde{N}_{xy} caused, for example, by finite filter overlap. $G_{noise} = c_x c_y T_e^2 \tilde{N}_{xy}^2$ is the portion of the cross-spectral density due to common noise between the two channels. As described below, this makes this derivation more general than most previous derivations that assume $\tilde{N}_{xy}^2 = 0$.

Similarly,

$$G_{xx}(f) = c_x^2 \left(\tilde{T}_e^2(f) + T_e^2 \tilde{N}_x^2(f) \right) \tag{C.7}$$

Note that this step differs from pervious derivations in two ways, both of which make the current derivation more generally applicable than previous derivations (for example, in Reference [105]). First, most previous derivations have assumed $\tilde{T}_e^2 \ll$

$T_e^2 \tilde{N}_x^2$, such that $G_{xx}(f) \approx c_x^2 (T_e^2 \tilde{N}_x^2)$ (note that the definitions of c_x and \tilde{N}_x^2 here are slightly different than those in Reference [105]). Relaxing this assumption leads to the eventual integrand of $\gamma_c/(1 - \gamma_c)$ in Equation 4.1 as opposed to just γ_c (aspect 2 from Section 4.2.1).

In addition, this derivation allows for some common noise between channels x and y by keeping a term G_{noise} , whereas previous derivations have assumed this term to be identically zero (this will later prove to be related to aspect 3 from Section 4.2.1). Keeping this additional term is an important extension over previous derivations.

Solving Equation C.7 for c_x gives:

$$c_x = \sqrt{\frac{G_{xx}(f)}{\tilde{T}_e^2(f) + T_e^2 \tilde{N}_f^2}} \quad (\text{C.8})$$

This assumes that the magnitude of the thermal noise on both channels is the same, even if the noise between the two channels is uncorrelated ($\tilde{N}_x^2(f) \approx \tilde{N}_y^2(f) \approx \tilde{N}_f^2$, where the f subscript emphasizes that this is noise per unit frequency). This assumption should be valid in any case where all CECE assumptions are valid (the two channels are closely enough spaced that they are sampling roughly the same background temperature). Such an assumption is also later justified by Equation C.16.

Now, solve Equation C.6 for $\tilde{T}_e^2(f)$:

$$\tilde{T}_e^2(f) = \frac{G_{xy}(f) - G_{noise}}{c_x c_y} \quad (\text{C.9})$$

Then insert Equation C.8 and the equivalent for c_y to get:

$$\tilde{T}_e^2(f) = \frac{G_{xy}(f) - G_{noise}}{\sqrt{G_{xx}(f)G_{yy}(f)}} \left(\tilde{T}_e^2(f) + T_e^2 \tilde{N}_f^2 \right) \quad (\text{C.10})$$

One can then define the complex coherence function, γ_c , (see page 390 of Ref.

[112]) as is done in Equation 4.2. For the sake of clarity it is repeated here:

$$\gamma_c(f) = \frac{G_{xy}(f)}{\sqrt{G_{xx}(f)G_{yy}(f)}} \quad (\text{C.11})$$

In addition, Equation 4.4 will be written in terms of G_{noise} , G_{xx} , and G_{yy} :

$$\gamma_{bg} = \frac{\overline{G_{noise}}}{\sqrt{\overline{G_{xx}}\overline{G_{yy}}}} \quad (\text{C.12})$$

where the overbar represents a frequency average over a frequency range far above the turbulent signal (assuming a constant background over the frequency range of interest, which is the case in all of the experimental discharges considered in this study). The value of γ_{bg} is calculated experimentally using Equation 4.4.

Using these definitions, one can write:

$$\tilde{T}_e^2(f) = (\gamma_c(f) - \gamma_{bg}) \left(\tilde{T}_e^2(f) + T_e^2 \tilde{N}_f^2 \right) \quad (\text{C.13})$$

Rearranging, solving for $\tilde{T}_e^2(f)$, and dividing by T_e^2 gives:

$$\frac{\tilde{T}_e^2(f)}{T_e^2} = \tilde{N}_f^2 \frac{\gamma_c(f) - \gamma_{bg}}{1 - (\gamma_c(f) - \gamma_{bg})} \quad (\text{C.14})$$

Integrating over the frequency range of interest and considering the root-mean-square value of the temperature fluctuation level (as in Reference [71] and others) gives:

$$\left. \frac{\tilde{T}}{\overline{T}} \right|_{rms} = \sqrt{\int_{f_1}^{f_2} \tilde{N}_f^2 \frac{Re\{\gamma_c(f) - \gamma_{bg}\}}{1 - Re\{\gamma_c(f) - \gamma_{bg}\}} df} \quad (\text{C.15})$$

where one takes the real part of the complex coherence function in order to obtain the root-mean-square value of the temperature fluctuations [115]. Note that this integral gives the band-limited turbulent temperature fluctuations from frequency f_1 to f_2 , which, as stated in the main text, are determined by the frequency width of the

turbulent feature.

For a finite frequency bandwidth (f_1 to f_2) one can substitute an expression for the thermal noise. From References [118, 120, 179], for an ECE radiometer the measured fluctuation level due to thermal noise is:

$$\frac{\tilde{T}_e^2}{T_e^2} \Big|_{noise} = \int_{f_1}^{f_2} \tilde{N}_f^2 df = B_{sig} \tilde{N}_f^2 = \frac{2B_{sig}}{B_{IF}} \quad (C.16)$$

where B_{IF} is the intermediate frequency bandwidth and $B_{sig} = f_2 - f_1$ is the signal bandwidth (allowing the signal bandwidth B_{sig} to be less than the video bandwidth). Thus, $\tilde{N}_f^2 = 2/B_{IF}$. If possible, it is also advantageous to measure \tilde{N}_f^2 experimentally, rather than using the theoretical value.

Inserting this into Equation C.15 gives the desired result, Equation 4.1:

$$\boxed{\frac{\tilde{T}}{T} \Big|_{rms} = \sqrt{\frac{2}{B_{IF}} \int_{f_1}^{f_2} \frac{Re\{\gamma_c(f) - \gamma_{bg}\}}{1 - Re\{\gamma_c(f) - \gamma_{bg}\}} df}} \quad (C.17)$$

Note that in Equation 4.1 the $\Big|_{rms}$ notation has been dropped for the sake of brevity.

C.2 Time Domain Derivation of Total Temperature Fluctuation

Most previous derivations of an equation for the total temperature fluctuation level from a CECE diagnostic have proceeded in the time domain, rather than the frequency domain [71, 105, 120]. Though the first derivation presented in this appendix was performed in the frequency domain, as this is clearer and more concise, it is also possible to perform the same derivation, starting with Equation C.1, in a manner more similar to past work. Parts of this derivation were also published in Reference [81].

Instead of Fourier transforming immediately, one defines the autocorrelation and

crosscorrelation functions as [115]:

$$R_{ij}(\tau) = \frac{1}{t_{tot}} \int_0^{t_{tot}} i(t)j(t + \tau)dt \quad (\text{C.18})$$

where t_{tot} is the total time period over which the measurement is taken, and $i(t)$ and $j(t)$ are two CECE signals.

Inserting the time domain signals into this definition, and taking only the fluctuating part at zero time delay gives equations that are the time domain analogues of Equations C.6 and C.7:

$$R_{xx}(0) = c_x^2 \left(\langle \tilde{\mathcal{T}}_e^2 \rangle + T_e^2 \langle \tilde{\mathcal{N}}_x^2 \rangle \right) \quad (\text{C.19})$$

and

$$R_{xy}(0) = c_x c_y \left(\langle \tilde{\mathcal{T}}_e^2 \rangle + T_e^2 \langle \tilde{\mathcal{N}}_{xy}^2 \rangle \right) = c_x c_y \langle \tilde{\mathcal{T}}_e^2 \rangle + R_{noise} \quad (\text{C.20})$$

where $R_{noise} = c_x c_y T_e^2 \langle \tilde{\mathcal{N}}_{xy}^2 \rangle$ is the portion of the cross-correlation function at zero time lag due to common noise between the two channels, the time domain equivalent of G_{noise} .

As before, one solves for c_x and c_y , combines equations, and solves for $\tilde{\mathcal{T}}$ to get:

$$\tilde{\mathcal{T}}_e^2 \approx \frac{R_{xy}(0) - R_{noise}}{\sqrt{R_{xx}(0)R_{yy}(0)}} (\tilde{\mathcal{T}}_e^2 + T_e^2 \tilde{\mathcal{N}}^2) \quad (\text{C.21})$$

Again, this assumes that the magnitude of the noise on both channels is the same, as in the first derivation.

Normalizing by T_e^2 , rearranging, and solving for $\tilde{\mathcal{T}}_e^2/T_e^2$ gives the analogue of

Equation C.10:

$$\frac{\tilde{T}_e^2}{T_e^2} = \tilde{\mathcal{N}}^2 \frac{\frac{R_{xy}(0)}{\sqrt{R_{xx}(0)R_{yy}(0)}} - \frac{R_{noise}}{\sqrt{R_{xx}(0)R_{yy}(0)}}}{1 - \left(\frac{R_{xy}(0)}{\sqrt{R_{xx}(0)R_{yy}(0)}} - \frac{R_{noise}}{\sqrt{R_{xx}(0)R_{yy}(0)}} \right)} \quad (\text{C.22})$$

Note that the $\tilde{\mathcal{N}}^2$ in this equation is different than the \tilde{N}_f^2 from Equation C.10 (the former is a total noise level while the latter is the noise level per unit frequency).

One then uses the relationship between correlation functions and spectral density functions in general (they are Fourier transforms on one another [115]), and specifically the relationships [180]:

$$\frac{R_{xy}(0)}{\sqrt{R_{xx}(0)R_{yy}(0)}} = \frac{1}{B_{sig}} \int_{f_1}^{f_2} Re\{\gamma_c(f)\} df \quad (\text{C.23})$$

and

$$\frac{R_{noise}}{\sqrt{R_{xx}(0)R_{yy}(0)}} = \frac{1}{B_{sig}} \int_{f_1}^{f_2} Re\{\gamma_{bg}\} df \quad (\text{C.24})$$

where B_{sig} and γ_c are defined as above. This relationship is exactly true in the limit of a continuous integral [180].

Taking the Fourier transform of Equation C.22 therefore gives:

$$\frac{\tilde{T}_e^2}{T_e^2} = \frac{\tilde{\mathcal{N}}^2}{B_{sig}} \int_{f_1}^{f_2} \frac{Re\{\gamma_c(f) - \gamma_{bg}\}}{1 - Re\{\gamma_c(f) - \gamma_{bg}\}} df \quad (\text{C.25})$$

Finally, using $\tilde{\mathcal{N}}^2 = B_{sig}\tilde{N}_f^2 = 2B_{sig}/B_{IF}$ and taking the square root gives the desired result, which is the same as for the frequency domain derivation (the same as Equation C.17):

$$\frac{\tilde{T}}{T} = \sqrt{\frac{2}{B_{IF}} \int_{f_1}^{f_2} \frac{Re\{\gamma_c(f) - \gamma_{bg}\}}{1 - Re\{\gamma_c(f) - \gamma_{bg}\}} df} \quad (\text{C.26})$$

C.3 Derivation of Uncertainty

This section of the appendix derives Equation 4.8, which is used to calculate the uncertainty in the total temperature fluctuation level from Equation 4.1.

Start with Equation 4.8:

$$\frac{\tilde{T}}{T} = \sqrt{\frac{2}{B_{IF}} \int_{f_1}^{f_2} \frac{Re\{\gamma_c(f) - \gamma_{bg}\}}{1 - Re\{\gamma_c(f) - \gamma_{bg}\}} df} \quad (C.27)$$

Now, assume that $Re\{\gamma_c(f) - \gamma_{bg}\} \ll 1$ such that:

$$\frac{\tilde{T}}{T} \approx \sqrt{\frac{2}{B_{IF}} \int_{f_1}^{f_2} Re\{\gamma_c(f) - \gamma_{bg}\} df} \quad (C.28)$$

Convert the continuous integral to a finite sum, as this is how the calculation is performed in practice:

$$\begin{aligned} \frac{\tilde{T}}{T} &\approx \sqrt{\frac{2}{B_{IF}} \sum_i \gamma_i \delta f} \\ &\approx \sqrt{\frac{2\delta f}{B_{IF}}} \cdot \sqrt{\gamma_1 + \gamma_2 + \gamma_3 + \dots} \end{aligned} \quad (C.29)$$

where δf is the frequency bin width for the finite sum calculation, and γ_i is the complex coherence function calculated for that bin.

Now, take the derivative of \tilde{T}/T with respect to γ_i and rearrange:

$$\begin{aligned} \frac{d}{d\gamma_i} \left(\frac{\tilde{T}}{T} \right) &\approx \sqrt{\frac{2\delta f}{B_{IF}}} \cdot \frac{1}{2\sqrt{\gamma_1 + \gamma_2 + \gamma_3 + \dots}} \\ &\approx \frac{\delta f}{B_{IF}} \frac{1}{\sqrt{2\delta f/B_{IF}} \sqrt{\gamma_1 + \gamma_2 + \gamma_3 + \dots}} \\ &\approx \frac{\delta f}{B_{IF}} \frac{1}{\tilde{T}/T} \end{aligned} \quad (C.30)$$

To calculate the total uncertainty on the calculation of \tilde{T}/T , one uses relatively standard error propagation techniques, writing:

$$\sigma_{\tilde{T}/T} = \sqrt{\sum_i \sigma_{\gamma_i}^2 \left(d/d\gamma_i \left(\frac{\tilde{T}}{T} \right) \right)^2} \quad (\text{C.31})$$

where σ_{γ_i} is the uncertainty on the complex coherence function calculated for each bin.

One then plugs Equation C.30 into Equation C.31 and rearranges to get:

$$\begin{aligned} \sigma_{\tilde{T}/T} &\approx \sqrt{\sum_i \sigma_{\gamma_i}^2 \left(\frac{\delta f}{B_{IF}} \frac{1}{\tilde{T}/T} \right)^2} \\ &\approx \frac{1}{\tilde{T}/T} \frac{1}{B_{IF}} \sqrt{\sum_i (\sigma_{\gamma_i} \delta f)^2} \end{aligned} \quad (\text{C.32})$$

which is the desired result, Equation 4.8.

C.4 Derivation of Sensitivity Limit

This section derives the sensitivity limit of the CECE diagnostic, which is given in Equation 4.9. This calculation involves integrating the uncertainty in the complex coherence function over the signal bandwidth.

To begin, start with Equation 4.1, assume $Re\{\gamma_c(f) - \gamma_{bg}\} \ll 1$ (which is consistent with trying to calculate the minimum detectable signal), and instead of integrating the complex coherence function γ_c , integrate its uncertainty σ_{γ_c} :

$$\left. \frac{\tilde{T}}{T} \right|_{limit}^{stat} > \sqrt{\frac{2}{B_{IF}} \int_{f_1}^{f_2} \sigma_{\gamma_c} df} \quad (\text{C.33})$$

An expression for σ_{γ_c} is given by Equation 4.7, and is repeated here for conve-

nience:

$$\sigma_{\gamma_c}(f) = \sqrt{\frac{1}{2n_d}(1 - |\gamma_c(f)|^2)^2} \quad (\text{C.34})$$

where n_d is the number of independent ensemble averaging windows. For the purposes of this derivation, assume $|\gamma_c(f)|^2 \ll 1$ (again, consistent with deriving a minimum detectable signal), which means $\sigma_{\gamma_c}(f) \approx \sqrt{1/2n_d}$. Inserting this into Equation C.33 gives:

$$\begin{aligned} \frac{\tilde{T}}{\bar{T}} \Big|_{limit}^{stat} &> \sqrt{\frac{2}{B_{IF}} \int_{f_1}^{f_2} \sqrt{\frac{1}{2n_d}} df} \\ &> \sqrt{\frac{2}{B_{IF}} \sqrt{\frac{1}{2n_d}} \int_{f_1}^{f_2} df} \\ &> \sqrt{\frac{2}{B_{IF}} \sqrt{\frac{1}{2n_d}} B_{sig}} \end{aligned} \quad (\text{C.35})$$

where $B_{sig} = f_2 - f_1$ is the signal bandwidth. Now, use the relationship between n_d and the total number of sample points, $n_d = NB_{sig}/B_{samp}$, where N is the total number of sample points and B_{samp} is the sampling rate. Inserting this into Equation C.35 to get:

$$\begin{aligned} \frac{\tilde{T}}{\bar{T}} \Big|_{limit}^{stat} &> \sqrt{\frac{2}{B_{IF}} \sqrt{\frac{B_{samp}}{2NB_{sig}}} B_{sig}} \\ &> \sqrt{2 \frac{B_{sig}}{B_{IF}} \sqrt{\frac{B_{samp}}{2NB_{sig}}}} \\ &> \sqrt{\frac{2}{\sqrt{N}} \frac{B_{sig}}{B_{IF}} \sqrt{\frac{B_{samp}}{2 \cdot B_{sig}}}} \end{aligned} \quad (\text{C.36})$$

which is the desired result, Equation 4.9.

In particular, if one takes uses the Nyquist Frequency as the signal bandwidth, $B_{sig} = B_{samp}/2$, then one finds that:

$$\frac{\tilde{T}}{\bar{T}} \Big|_{limit}^{rad} > \sqrt{\frac{2}{\sqrt{N}} \frac{B_{sig}}{B_{IF}}} \quad (C.37)$$

which is Equation 4.10, and is the formula used in References [71, 120]. This difference between Equations 4.9 and 4.10 is consistent with suggestions in References [73, 121] that one must correct the original formula for the sensitivity limit when the signal bandwidth and sampling rate differ significantly.

This completes the derivation of the CECE sensitivity limit.

C.5 Optical Depth Derivation

This final section of the appendix derives Equation 2.29, which is used to calculate the optical depth of second harmonic X-mode electron cyclotron emission. This optical depth calculation is used to ensure that the CECE diagnostic is truly measuring electron temperature fluctuations, and not some combination of density and temperature fluctuations.

This derivation does not go through the entire process of obtaining the optical depth, instead continuing the derivation already presented in Reference [68] in order to achieve the final equation used in practical calculations. This derivation therefore begins with the equation from the third row of Table XII from Reference [68], giving the generic expression for $n \geq 2$ harmonics of both modes of propagation:

$$\tau_n^{(O,X)} = \frac{\pi^2 n^{2(n-1)}}{2^{n-1} (n-1)!} \left(\frac{\omega_{pe}}{\omega_{ce}} \right)^2 \left(\frac{v_{Te}}{c} \right)^{2(n-1)} (\sin \theta)^{2(n-1)} (1 + \cos^2 \theta) \mu_n^{(O,X)}(\theta) \frac{L_B}{\lambda_O} \quad (C.38)$$

where n is the emission harmonic, ω_{pe} is the electron plasma frequency, ω_{ce} is the electron cyclotron frequency, v_{Te} is the electron thermal velocity, c is the speed of

light, θ is the angle between the propagation and the background magnetic field, μ_n is a function that describes polarization effects in the wave absorption, $L_B = B_0|dB_0/ds|$ (s is the path length traversed by the radiation), and $\lambda_O = 2\pi c/\omega_{ce}$.

From Equation 3.1.79b [68]:

$$\mu_n^X(\theta) = \left(N'_{\perp,X}\right)^{2n-3} \left[1 + \frac{(\omega_{pe}/\omega_{ce})^2}{n[n^2 - 1 - (\omega_{pe}/\omega_{ce})^2]}\right]^2 \quad (\text{C.39})$$

where $N'_{\perp,X}$ is the X-mode refractive index, and is given by Equation 3.1.12 [68]:

$$(N'_{\perp})^2 = 1 - \left(\frac{\omega_{pe}}{\omega}\right)^2 \frac{\omega^2 - \omega_{pe}^2}{\omega^2 - \omega_{ce}^2 - \omega_{pe}^2} \quad (\text{C.40})$$

Now, define:

$$q = \left(\frac{\omega_{pe}}{\omega_{ce}}\right)^2 \quad (\text{C.41})$$

Letting $\omega = 2\omega_{ce}$ in Equation C.40 and using the definition of q gives:

$$(N'_{\perp})^2 = \frac{12 - 8q + q^2}{12 - 4q} \quad (\text{C.42})$$

Inserting this into Equation C.39 and using $n = 2$ for the second harmonic one obtains:

$$\mu_n^X(\theta) = \left(\frac{12 - 8q + q^2}{12 - 4q}\right)^{1/2} \left(\frac{6 - q}{6 - 2q}\right)^2 \quad (\text{C.43})$$

Finally, insert this into Equation C.38, assume perpendicular propagation such that $\theta = \pi/2$, again use $n = 2$ for the second harmonic, use $v_{Te}^2 = T_e/m_{e0}$ (defined in this derivation as used in Reference [68], without the standard factor of 2), and use $\lambda_O = 2\pi c/\omega_{ce}$ to obtain the desired result:

$$\tau_2^X = \pi q \frac{T_e}{m_{e0}c^2} \frac{\omega_{ce}}{c} \left(\frac{12 - 8q + q^2}{12 - 4q}\right)^{1/2} \left(\frac{6 - q}{6 - 2q}\right)^2 L_B \quad (\text{C.44})$$

This is Equation 2.29, and is the same result shown as Equation 4.53 in Reference [67].

Appendix D

The Ricci Validation Metric

This appendix contains more information about the Ricci validation metric that was used in Chapters 5 and 6 of this thesis. The metric itself is originally from Reference [149]. This appendix includes the equations required to calculate the metric, as well as the weighting choices that were used in this thesis work.

The metric itself, χ_{Ricci} , is defined as:

$$\chi_{Ricci} = \frac{\sum_j R_j H_j S_j}{\sum_j H_j S_j} \quad (D.1)$$

where the sum over j is the sum over each validation constraint. R_j quantifies the level of agreement between the experimental measurement and the simulation output, H_j denotes the level in the validation hierarchy for a given constraint, and S_j quantifies the uncertainties in the simulation and experimental value for a given constraint. In this thesis, the four primary constraints are the electron heat flux, the ion heat flux, the perturbative thermal diffusivity, and the total electron temperature fluctuation level. The definitions of each of these inputs is now explained in greater detail.

The level of agreement between experiment and simulation, R_j , is defined as:

$$R_j = \frac{\tanh [(d_j - d_0)/\lambda] + 1}{2} \quad (\text{D.2})$$

where \tanh is the hyperbolic tangent function, d_j is the normalized distance between experiment and simulation, d_0 is the value of d_j that dictates the transition from agreement to disagreement, and λ dictates the sharpness of the transition. As is recommended in Reference [149], this thesis uses $d_0 = 1.5$ and $\lambda = 0.5$ for all calculations of χ_{Ricci} . Others using the Ricci Metric may make other choices for these parameters, so it is important to state what choices are made before describing one's results.

The normalized distance for each constraint, d_j , is defined as:

$$d_j = \sqrt{\frac{1}{N_j} \sum_{i=1}^{N_j} \frac{(x_{j,i} - y_{j,i})^2}{\Delta x_{j,i}^2 + \Delta y_{j,i}^2}} \quad (\text{D.3})$$

where $x_{j,i}$ is the experimental measurement of constraint j at point i (allowing for comparisons of several points of the same constraint, for example two frequency bins of a measured frequency spectrum), $y_{j,i}$ is the simulation value of constraint j at point i , $\Delta x_{j,i}$ is the uncertainty of the experimental measurement, and $\Delta y_{j,i}$ is the uncertainty of the simulation value. In this thesis, all validation constraints are a single value, so there is no sum over i .

Turning now the hierarchy level, H_j , this thesis doesn't make any assumptions or judgements about the hierarchy difference between different constraints, and so sets $H_j = 1$ for every validation constraint. Please see Reference [49] for further discussion of the validation hierarchy and how one might make judgements about which constraints are appropriate on higher or lower levels of the hierarchy.

Finally, the uncertainty of each constraint is quantified as:

$$S_j = \exp\left(-\frac{\sum_i \Delta x_{j,i} + \sum_i \Delta y_{j,i}}{\sum_i |x_{j,i}| + \sum_i |y_{j,i}|}\right) \quad (\text{D.4})$$

where all quantities are defined above. In the case of zero uncertainty, $S_j = 1$, and as the uncertainty for a constraint increases, the its weight in the final metric decreases. In other words, S_j gives more weight to constraints that are measured with lower uncertainty, and less weight to constraints that are measured with higher uncertainty.

This completes the equations required to calculate the Ricci Metric, χ_{Ricci} . As shown in this appendix, the Ricci Metric gives a weighted sum of the discrepancy between experimental and simulation values of several validation constraints. Those constraints with smaller uncertainties are weighted more heavily than those with larger uncertainties. The metric also allows one to weight different levels of the validation hierarchy differently, but this thesis did not do so. The three parameters that one chooses when calculating the Ricci Metric are H_j (the hierarchy weight of each constraint), d_0 (the transition value from agreement to disagreement for a given constraint), and λ (the sharpness of the cutoff transition), for which this thesis used $H_j = 1$, $d_0 = 1.5$, and $\lambda = 0.5$, which are the values recommended in Reference [149].

Appendix E

Density Fluctuation Measurements with Reflectometry on the Large Helical Device (LHD)

This appendix describes the development of new analysis techniques to calculate radial profiles of density fluctuations measured with reflectometry on the Large Helical Device (LHD). This work was completed as part of this thesis and was originally published in Reference [181], and though it does concern turbulence measurements, it did not fit directly into the narrative of the thesis and is therefore described in this appendix.

In this work, fast far infrared laser interferometry (FIR) density profile data was combined with lower time resolution results of the ray tracing code LHDGauss in order to determine the cutoff location of a reflectometer diagnostic with high time resolution. The plasma density was quickly scanned by injecting a pellet, leading to a scan in the reflectometer measurement location. The reflectometer data was used to locally measure both plasma rotation velocity and density fluctuations. Rotation velocity profile measurements agreed with measurements made by charge exchange

spectroscopy (CXS), once diagnostic differences were accounted for. The local density fluctuation measurements showed strong peaking near the edge of the plasma, with a full width at half maximum of 1.5 - 3 cm (less than 5% of the normalized minor radius), which is consistent with data from other machines. These measurements may also be used in the future to improve inversion algorithms for phase contrast imaging (PCI) diagnostics.

The remainder of this appendix closely follows Reference [181].

E.1 Introduction to Density Fluctuation Measurements

In addition to the temperature fluctuations used throughout this thesis, measuring density fluctuations in fusion plasmas is also key to better understanding and modeling turbulence and turbulent transport. Just as temperature fluctuations can be measured with the CECE diagnostic described in Chapter 4, a number of diagnostics can be used to measure density fluctuations, including: beam emission spectroscopy (BES) [182], phase contrast imaging (PCI) [183], and reflectometry [184, 185, 186]. Fortunately, most of these techniques are not constrained by the optical depth near the edge of plasmas as CECE measurements are, and density fluctuations have been observed to increase sharply near the edge of tokamak plasmas on a number of machines worldwide [187, 188, 189, 190, 191, 192].

As described in Chapter 2, the edge and pedestal region is often qualitatively different than the core plasma region, which has been the main focus of thesis. Nevertheless, turbulence in this region is governed by the same general principles as it is in the core region, and so the increase in density fluctuations near the edge is likely related to turbulence drive and suppression terms. First, normalized gradient scale lengths are typically very large near the edge, leading to increased turbulent drive [193]. In addition, there is a very sharp increase in the safety factor q profile, which

describes the angle of the magnetic field in the torus. It is possible that this strongly changing magnetic configuration, in particular the magnetic shear, also plays a role in the measured increase in fluctuations.

Historically, most edge density measurements have been made on tokamaks, partially due to their greater number worldwide. In order to investigate the role of magnetic geometry in driving edge density fluctuations, however, it is important to measure this phenomenon in other magnetic configurations, such as helical devices. Very briefly, stellarator-type devices generate a stable plasma through carefully shaped magnetic field coils, rather than by driving a current in the plasma like a tokamak. While there are many examples of single location measurements, only a few fluctuation profile measurements on stellarators have been published: Reference [194] on ATF, References [195, 196] on TJ-II, and Reference [197] on Wendelstein 7AS. The work described here makes these measurements on the Large Helical Device (LHD) [198, 199] using reflectometry.

There have been some measurements of density fluctuation profiles in the past on LHD with PCI [200, 201, 202, 203], but these are all based on line-integrated measurements. For this reason, the profile must be inverted from measurements along several chords. The inversion process is mathematically complicated, and may lead to large uncertainty in the inferred profile. Specifically, this past work has shown that the inferred profile depended heavily on the inversion method, with some methods leading to profiles that were peaked fairly far in radially. In addition, PCI diagnostics are susceptible to systematic errors due to optical misalignment, which may affect the measured location of the peak of the fluctuation profile [202].

The work described in this appendix attempts to overcome these difficulties by developing new analysis techniques enabling local density fluctuation profile measurements on LHD with reflectometry. Reflectometers measure density fluctuations at only one location at a given time, and so profile measurements are accomplished

by scanning the reflectometer cutoff with a pellet-induced density scan. The new analysis technique focuses on fast mapping of the reflectometry measurement location during this scan, and enables local measurement of density fluctuation profiles. These profiles may also be used to inform inversion and systematic error correction of PCI measurements.

The remainder of this appendix is organized as follows. Section E.2 describes the diagnostic hardware used in this study. Section E.3 describes the new analysis methods developed as part of this work. Finally, Section E.4 presents profile measurements of both rotation velocity and density fluctuations, as well as evidence that these the profiles are not affected by temporal variation in the density fluctuation levels caused by the pellet injection.

E.2 Experimental Set-Up on LHD

Though a detailed description is not given here, Doppler reflectometry (or Doppler back-scattering) is a relatively widespread technique used to measure turbulent flow velocity and amplitude in fusion plasmas. References [184, 185, 186] contain considerable detail on the physical principles behind this diagnostic as well as how it is implemented in practice. In essence, a reflectometer launches a wave into the plasma that reflects at a particular electron density. If this wave is launched at an angle to the plasma’s density gradient, the Doppler shift of the reflected waves gives information about the plasma rotation velocity and density fluctuation amplitude.

Specifically, two Doppler reflectometers have been installed on LHD: a frequency hopping system [204] and a multi-channel frequency comb system [205]. For historical reasons, the frequency-hopping system is referred to as the ‘9-O’ reflectometer, since it is installed on the ‘9-O’ toroidal port, and the multi-channel system is referred to as the ‘comb’ reflectometer (it is installed on the ‘3-O’ toroidal port). The 9-O system is typically set to launch waves at 30.0 GHz, and the comb system operates

with frequencies of 27.7, 29.0, 30.5, 32.0, 33.3, 34.8, 37.0, and 38.3 GHz. These reflectometers measure turbulent electron density fluctuation perpendicular wavenumbers up to approximately 15 cm^{-1} , with a resolution of approximately 2 cm^{-1} [186, 205]. Both systems use vertically separated bi-static antennae and generally operate in the O-mode polarization. The channels of the comb reflectometer are not calibrated to one another, and each measures density fluctuations in their own arbitrary units.

While there are many difficulties in interpreting Doppler reflectometry data, the one that is the focus of this work is estimating the reflection location, and thus the location at which the reflectometer is measuring density fluctuations. In the simplest treatment of an O-mode reflectometer, one uses cold plasma equations to calculate the density at which the launched wave frequency is equal to the plasma frequency. In this case, one only requires knowledge of the plasma density profile. Unfortunately, the complexities of wave propagation in hot plasmas often necessitate ray tracing codes in order to accurately assess the reflection location [206, 207].

As their name suggests, ray tracing codes trace a ray of the launched wave as it propagates through the plasma, using a variety of models for the plasma index of refraction. The work presented here uses the 3D ray tracing code LHDGauss, which, in its most recent update, has been mainly applied to ECH experiments in LHD [206, 207]. This code solves the eikonal equation under the WKB approximation. The equilibrium electron density profile is provided by 3D equilibrium mapping of Thomson Scattering measurement [208, 209]. This work will also include density measurements from a far infrared laser interferometer (FIR) [210].

Figure E-1 shows an example of the output from LHDGauss. The reflectometer measurement location is determined by the point at which the wave reflects (the point of its maximum radial penetration). The figure also shows the cutoff location as calculated with a simple cold plasma reflection as the thin lines, revealing that the two calculations disagree quite significantly in some cases. The main uncertainties in these




Figure E-1: Example ray trajectories in a horizontal plasma cross section calculated with the ray tracing LHDGauss code. The thick solid lines indicate the rays of reflectometer probing beams at 27.7, 30.5 and 33.3 GHz. The thinner and dashed lines are the nominal cutoff densities, assuming cutoff when the plasma frequency is equal to the beam frequency. The ray tracing results show reflection before the nominal cutoff. Originally from Reference [181].

calculations arise from uncertainties in the plasma density profile. LHDGauss calculations have shown that uncertainties arising from reflectometer beam misalignments are negligible compared to other uncertainties in the analysis.

In addition to equilibrium profile measurements, Doppler reflectometry, and PCI data, this work also uses rotation velocity measurements from poloidal charge exchange spectroscopy (CXS) [211, 212] in order to compare to the reflectometry data and check the mapping. The CXS system has two poloidal views, one upward and one downward, and calculates the rotation velocity by using the difference in the Doppler shift observed on the two views. The effective spatial resolution is approximately 3 cm.

E.3 Reflectometry Analysis Methods

This section describes the new analysis procedure developed to extract the fluctuation profile from the reflectometer fluctuation data. Before diving into the details, a brief

summary is given.

Since each reflectometer channel measures the fluctuation level at only one location (and the channels of the comb reflectometer are not calibrated to one another), one must scan the reflectometer cutoff location in order to measure a fluctuation profile. This is achieved experimentally by injecting a pellet into the plasma. The pellet transiently increases the plasma density, which then transiently scans the reflectometer measurement location outward.

As discussed in the last section, the 3D ray tracing code LHDGauss is needed to determine the exact reflectometer measurement location. Running this code at every 1ms time slice required to resolve the fast cutoff scan caused by a pellet induced density scan tends to be prohibitively slow, and so a faster estimate of the cutoff location was required. This is achieved using a scaled cutoff density in the simple plasma frequency cutoff calculation. The cutoff density scaling is chosen to match the measurement location calculated with ray tracing for each Thomson Scattering slice (every 33 ms). Between these time slices, the scaled cutoff density is applied to the faster FIR density profiles. By calculating the reflectometer measurement location at each time step, one can map the time signal of the reflectometer fluctuations to a spatial location, thus mapping out a profile. The rotation velocity profile measured in this manner was then checked against that measured by CXS in order to confirm the validity of the analysis method.

Once the cutoff density scaling has been determined using ray tracing and checked against the velocity profile measured with CXS, the same technique can be used to measure the density fluctuation profile. One can also use the peak of the fluctuation profile to calibrate the channels of the comb reflectometer to the 9-O reflectometer. Once calibrated to one another, one can use all of the channels in order to measure fluctuation profiles even in periods of constant density.

Profile Fitting and Reflectometer Cutoff Calculation

The first step of the analysis procedure for the profile measurement is to fit the density profile and calculate the reflectometer measurement location at each time step. The two density profile diagnostics on LHD are the FIR and Thomson Scattering. The FIR has high time resolution (1 ms in its typical configuration), but low spatial resolution. Thomson Scattering, on the other hand, has high spatial resolution, but slow time resolution (33 ms). As mentioned above, this work used the fast FIR profiles, and checked them against the Thomson profiles as appropriate.




Figure E-2: The electron density profile of LHD discharge 116119, measured with both Thomson Scattering and FIR. The Thomson data is smoothed with a Savitzky-Golay filter, and the FIR data is spline fit. The nominal cutoff density for the 9-O reflectometer is marked as the horizontal black line. Near the cutoff density, the FIR and Thomson fits agree quite well. r_{eff} is the effective minor radius, if the plasma cross section was mapped to a circle. Originally from Reference [181].

Since the FIR is a line-integrated diagnostic, the first step in calculating the profile is a standard inversion of the line-integrated data [210]. A spline fit was used to fit the inverted profile, and example of which is shown in Figure E-2, along with a profile fit to Thomson data. While there is some discrepancy between the two profiles, most clearly around $r_{eff} = 0.5$ m, the two diagnostics agree quite well near the plasma

edge, which is the region most important in this study.

In order to check first the results of an overly simplified analysis, the reflectometer measurement location was calculated as the location where the density first increased enough such that the plasma frequency was equal to the reflectometer launch frequency. The time evolution of this cutoff location calculation after a pellet injection is shown as the red trace in Figure E-3. The other traces on this plot are addressed later in this section. Comparing the results of this simplified analysis with the ray tracing results in the next section will show the importance of developing the new technique in this work, based on scaling the cutoff density.




Figure E-3: The 9-O reflectometer measurement location after a pellet injection calculated with scaled cutoff densities and with ray tracing. The solid lines are calculated with FIR profiles and a scaled cutoff density. The magenta squares are calculated with Thomson Scattering data and a scaled cutoff density. The cyan circles are calculated with the Thomson Scattering data and ray tracing. A cutoff density scaling of 0.75 agrees most closely with the ray tracing results for both the FIR and Thomson Scattering data. Figure originally from Reference [181].

Scaling of Cutoff Density to Match Ray Tracing and CXS Velocity Profiles

Using the fit density profiles from the FIR diagnostic, one can compare the results of a scaled cutoff density to those of full ray tracing results in order to determine an

appropriate scaling. The accuracy of the resulting scaling is checked by comparing the rotation velocity profile measured in this manner to that measured with CXS.

Figure E-3 shows the reflectometer measurement location as calculated in a variety of ways. As discussed before, the red line is based on an unscaled cutoff density from the plasma frequency. The other colored lines use instead a scaled cutoff density, $n_{cutoff,s} = s \cdot n_{c,0}$, where $n_{c,0} = m_e \epsilon_0 \omega^2 / e^2$ and s is the scaling factor. Ray tracing is run using Thomson Scattering profiles every 33 ms, the results of which are represented by the cyan circles. As the figure shows, in this particular discharge, a cutoff density scaling of 0.75 achieves the best agreement with the ray tracing results. More generally, similar analysis on a number of other discharges showed that a scaling of between 0.7 and 0.8 agreed well with ray tracing. In addition, Figure E-3 shows the reflectometer measurement location calculated based on a scaled cutoff density and the Thomson Scattering density profile, for which a scaling of 0.75 also works well.

By mapping the reflectometer signal at each time to the measurement location at each time, one can therefore map out a spatial profile of density fluctuation and velocity measurements using the two time series data. As confirmation of the accuracy of the mapping of the reflectometer measurement location during the pellet induced density scan, the velocity profile measured in this way is compared to the velocity profile measured with CXS.

The velocity profiles mapped from a variety of cutoff scalings are shown in Figure E-4 alongside the velocity profile measured with CXS. Note that the comparison of reflectometer and CXS data is meant *only* to confirm the spatial mapping of the reflectometer data, and is not meant as a detailed comparison of velocity measurements in general between these two different diagnostics. Reference [184] contains further discussion of diagnostic comparison.

The comparison in Figure E-4 confirms that a scaling of 0.75 for the reflectometer cutoff density produces the best agreement with CXS in the shape and peak location




Figure E-4: Perpendicular velocity profiles measured with CXS and the 9-O reflectometer after a pellet-induced density scan. The gray profile is the CXS data. Different colors represent different cutoff density scalings for the reflectometer mapping. Diagnostic uncertainty is shown as the shaded regions. The closest agreement with CXS for the peak of the velocity profile is achieved with a scaling of 0.75. Figure originally from Reference [181].

of the velocity profile. The maximum velocity differs by approximately 30%, while the width of the profile peak is similar. The peaks of the profiles measured with the two diagnostics agree within uncertainty. The slight disagreement in the peak velocities does not alter the assessment that a scaling of 0.75 best maps the reflectometer data.

Figure E-4 also shows that the smaller r_{eff} side of the profile seems to agree more closely with the CXS data than the larger r_{eff} side. This is perhaps not surprising, as the reflectometer measurements are expected to be less reliable right at the plasma edge where density fluctuations are the largest.

Finally, it is of interest to estimate the uncertainty associated with the cutoff mapping procedure just described. This is done by varying the cutoff density scaling and checking agreement with ray tracing results within the ray tracing uncertainties. One then uses the different scalings to map the measured velocity profile, and compares these results to CXS measurements. Both of these comparisons give an uncertainty

of approximately ± 0.05 in the cutoff scaling. This uncertainty will be used in the calibration of the comb reflectometer later in this section and in the fluctuation profile mapping in Section E.4.

The results of the comparison of the density cutoff scaling with ray tracing results, and the mapped velocity profiles to velocity profiles measured with CXS, indicate that this method of obtaining profile measurements with reflectometry during a pellet induced density scan is suitable for use with density fluctuations.

Calibration of Comb Reflectometer

In addition to using the density scan to choose the appropriate cutoff density scaling for the reflectometer measurement location mapping, the scan can also be used to calibrate the channels of the comb reflectometer to one another and to the 9-O reflectometer. This process is worthwhile since the 9-O velocity measurements are absolutely calibrated to measure km/s, while the comb reflectometer measures in arbitrary units that are different for each channel. For this reason, without cross-calibration one cannot compare the channels to one another or to the 9-O reflectometer.

By calculating the mapped velocity profiles from each of the comb reflectometer channels and the 9-O reflectometer, and plotting them over one another, one can calibrate the comb channels such that they also measure the velocity in km/s. Figure E-5 shows the profiles measured with the comb and 9-O reflectometers after calibration. Three comb channels were not operated in this discharge, so are omitted from the figure. All six profiles agree quite well within uncertainty.

Once this calibration has been completed, one can compare the different channels of the comb reflectometer to each other during other times in the discharge. In this manner, one can measure profiles even in time periods of constant density. More importantly, a similar procedure can also be applied to density fluctuation profiles, for which there are no other local measurement against which to check.




Figure E-5: Perpendicular velocity profiles measured with calibrated data from five channels of the comb reflectometer. A cutoff density scaling of 0.75 was used. The peak of the velocity profile measured with each channel was calibrated to the 9-O reflectometer. Representative velocity (shaded region) and mapping (black error bars) uncertainties are shown for the 32.0 GHz channel. Figure originally from Reference [181].

E.4 Density Fluctuation Profile Measurements

Using the cutoff scaling procedure from the previous section, one can use the pellet-induced density scan to measure the density fluctuation profile in the same manner as the velocity profiles. After addressing a possible complication in the analysis method, this procedure is applied to LHD discharge 116119, the temperature and density profiles of which are shown in Figure E-6.

Separation of Fluctuation Changes Due to Pellet and Due to Density Scan

Before showing the measured density fluctuation profiles, one possible complication must first be addressed. Since the procedure outlined in this work takes advantage of a pellet-induced density scan to scan the reflectometer measurement location and measure the fluctuation profile, one must ensure that the pellet injection itself is not responsible for the change in fluctuations. In other words, one must make sure that




Figure E-6: Temperature and density profiles near the edge of LHD discharge 116119. Raw data is shown as symbols and fits are shown as lines. Figure originally from Reference [181].

the observed change in fluctuations is indeed a spatial variation, and not a temporal variation caused by the pellet injection.

This check is performed using PCI measurements of density fluctuations. Unlike the reflectometer, the PCI chords are fixed in place, and do not change their measurement location as the plasma density changes. The PCI therefore directly measures temporal changes in fluctuations caused by the pellet injection. By comparing the time histories of the PCI and reflectometer fluctuation measurements, one is therefore able to check whether or not the reflectometer fluctuation measurement changes are caused by the pellet itself, or by the scan in measurement location.

Again, as with the CXS velocity measurements, this comparison is *not* intended as a direct comparison of the two diagnostics, especially since the PCI and reflectometer diagnostics on LHD measure different wavenumber spectra, and is only intended as evidence that the fluctuation changes caused by the pellet injection are temporally separated from the reflectometer measurement location scan.

Figure E-7 shows time histories of fluctuations measured with one PCI channel




Figure E-7: Time history of density fluctuations measured with one line-integrated PCI chord (50-150 kHz) and the 9-O reflectometer (30-150 kHz) just after a pellet injection. The region highlighted in red shows changes in the fluctuation levels associated with the pellet injection itself. The region highlighted in blue shows the fluctuation changes in the reflectometer signal associated with the change in measurement location caused by the density scan. These features are clearly separated temporally, indicating that the second feature is not influenced by the pellet injection itself. Figure originally from Reference [181].

(line integrated, so measuring both edge and core, 50-150 kHz fluctuations) and the 9-O reflectometer (30-150 kHz). The pellet is injected at approximately 4.56 seconds, and causes a sharp increase in fluctuation levels measured by both PCI and reflectometry, which quickly decays as the pellet-induced fluctuations dissipate. Then, starting at approximately 4.58 seconds, the reflectometer begins to measure increased fluctuation levels while the PCI sees no change in fluctuation level. Since this feature is absent on the PCI, there is strong evidence that the feature is due to a change in the reflectometer measurement location, and not due to a temporal change in fluctuation levels.

With this comparison as evidence that the pellet injection itself is not responsible for the observed increase in fluctuation level, this work proceeds to map out density fluctuation profiles with the scanned reflectometer measurements.

Fluctuation Profile from 9-O Reflectometer

Figure E-8 shows density fluctuation profiles measured with mapped data from the 9-O reflectometer after a pellet injection. Three frequency ranges are shown, 3-30 kHz, 30-150 kHz, and 150-490 kHz. These frequency ranges are not adjusted for the Doppler shift, and are calculated as in Reference [204]. This figure shows that the density fluctuations are strongly peaked near the edge of the plasma in all three frequency ranges. Note that the lowest frequency range may contain some contributions from MHD activity, which is not the primary focus of this work.




Figure E-8: Density fluctuation profiles in three frequency ranges measured with the 9-O reflectometer during a pellet-induced density scan and a cutoff density scaling of 0.75. The lowest frequency range may contain some contributions from MHD activity. The vertical dotted line is a_{99} , which is defined as the effective minor radius within which 99% of the total plasma stored energy is confined. Fluctuations are strongly peaked near the plasma edge. Fluctuation level uncertainty is shown as the shaded regions, and radial mapping uncertainty as the error bars. Figure originally from Reference [181].

Returning to one of the motivations for this mapping procedure, the profiles shown in Figure E-8 agree roughly with some of the PCI inversions in Reference [203] (for example the maximum entropy inversion method), which show a strong peak near the edge of the plasma, but disagree with other inversions (such as the auto-regressive

Discharge	Fluctuation FWHM	Fluctuation FWHM
	30-150 kHz (cm) ± 0.25 cm	150-490 kHz (cm) ± 0.25 cm
116119	1.81	1.52
119128	2.32	1.41
119129	2.02	1.55
119130	2.09	1.25
119131	2.04	1.09
119132	2.37	1.53
119133	1.07	1.38

Table E.1: The full width at half-maximum of the edge density fluctuation peak measured in a series of discharges in LHD with the 9-O reflectometer after a pellet injection. Table originally from Reference [181].

method). The profile shapes also generally agree with what is seen on tokamaks. While PCI inversions have not yet been completed on this discharge, the reflectometer data presented here will inform future inversions of PCI data.

Note that these two diagnostics measure different turbulent wavenumbers (reflectometry measures a specific k up to 15 cm^{-1} [186], and PCI measures an integrated k spectrum from 1 to 9 cm^{-1} [203]). For this reason, and as discussed earlier in this appendix, a quantitative comparison of turbulence measurements from these two diagnostics is inappropriate, but qualitative comparisons of profile shapes are still worthwhile.

In the discharge discussed here, and in others analyzed as part of this work, the radial width of the density fluctuation feature seems to be the largest in the 30 -150 kHz (with a full width at half-maximum of approximately 1.5 - 3 cm), and decreases at higher frequencies. A summary of the radial widths of the density fluctuation profile from a series of discharges is given in Table E.1. These discharges varied only slightly

in plasma parameters, and their analysis was intended primarily to verify that the procedure developed in this work could be applied to a number of other discharges. Parametric dependencies of the fluctuation profile width may be the subject of future work.

Fluctuation Profile from Comb Reflectometer

Finally, in addition to using data from the 9-O reflectometer, the edge density fluctuation profile was also measured with the various channels of the comb reflectometer. As was performed for velocity measurements, the fluctuation measurements from the comb reflectometer must be calibrated using the peak of the fluctuation profile from the 9-O reflectometer. Though the final data still has arbitrary units, the channels will at least all be measuring in the same arbitrary units.




Figure E-9: Density fluctuation profiles measured with five channels of the comb reflectometer. A cutoff density scaling of 0.75 was used, and the amplitude of each channel was calibrated to the 9-O measurement. The vertical dotted line is a_{99} . Representative uncertainties are shown for the 32.0 GHz channel, with fluctuation level uncertainty as the shaded region and radial mapping uncertainty as the error bars. Figure originally from Reference [181].

Figure E-9 shows the calibrated comb reflectometer fluctuation measurements.

While the agreement across channels is not as clean as that for the velocity profile, shown in Figure E-5, the channels show a very similar fluctuation profile shape and peak location. There is some discrepancy, especially in the inner radii, but this is likely related to the noise floor of the comb reflectometer channels, which is larger than for the 9-O reflectometer.

Conclusions

This appendix has described a new method of measuring density fluctuation profiles on LHD using Doppler reflectometry during a pellet-induced density scan. The quickly changing density profile after a pellet injection is measured with 1ms time resolution using the FIR diagnostic. The reflectometer measurement location is calculated using a scaled cutoff density, matched to ray tracing results from the LHDGauss code. By mapping the reflectometer measurement to a radial location at each time, one is able to measure radial profiles of rotation velocity and density fluctuations. The measured velocity profiles agree qualitatively with those measured with CXS.

The density fluctuation profiles measured in this manner peak strongly near the edge of the plasma, which is consistent with past work on both tokamaks [187, 188, 189, 190, 191, 192] and other helical devices [195, 196, 197]. The full width at half maximum of this fluctuation peak is approximately 1.5 - 3 cm for the LHD discharges analyzed here, less than 5% of the normalized minor radius. These results agree with the results of some PCI inversion methods, but disagree with others, providing further evidence as to which methods are accurately capturing the fluctuation profile in LHD.

In addition to informing which inversion methods are appropriate for PCI data, the density fluctuation profiles measured with this analysis technique may in the future be used as a validation constraint for turbulent transport models, in the same way that the electron temperature fluctuations from CECE were used in Chapters 5 and 6 in this thesis. While the specific measurements made in this appendix were

very near the edge of the plasma, in theory the same technique could be used deeper in the core of the plasma, adding another constraint to the set already used in this thesis.

Appendix F

MQ1 Tokamak Design using the Tokamak Simulation Code (TSC)

One of the original motivations of the work in this thesis was the prediction of future machine performance. Ideally, one would like to be able input a machine design (including magnetic coil locations and currents, heating systems, fueling, etc.) and then simulate an entire plasma discharge, from breakdown to peak performance to ramp down. The work in this thesis primarily focused on validating the models for turbulent heat transport in the core of the tokamak plasma, but to simulate the entire tokamak, one requires models for every other aspect of the machine and plasma as well, including the macroscopic magnetohydrodynamic evolution, the edge and pedestal transport, heating systems and others.

While performing such simulations with high physics fidelity is beyond the capabilities of the fusion community at the moment (both due to limited computational power and limited fundamental understanding of the physics behind, for example, pedestal formation in H-mode), it is possible to perform such machine simulations using highly simplified models. In fact, this is how most machines are designed. At the most basic level, one wants to design a set of magnetic coils that will successfully

generate and confine a plasma. Since the plasma itself has a strong impact on the dynamics of the full system (especially on the ability to drive plasma current), it is essential to include at least some model for plasma transport in such a simulation.

One example of such a design code is the Tokamak Simulation Code (TSC) [213]. The primary use of this code is to design the magnetic coils for tokamaks. To calculate the macroscopic stability and evolution of the plasma, and therefore whether or not a given set of coils and heating will properly confine a plasma, TSC solves the Grad-Shafranov equation, which describes time-dependent MHD in a toroidal geometry [6]. Doing so requires information about the plasma temperature, density, and current profiles. Since TSC was developed before the advent of modern turbulent transport models and modern supercomputers, and really only aims to have accurate enough plasma profiles to inform coil design, it uses highly simplified analytic models of heat and particle transport. Specifically, it uses the Coppi-Tang model [214, 215] to model core transport, a simplified model of the sawtooth crash, and approximates the H-mode profile by artificially reducing transport in a narrow region near the plasma edge.

While these approximations are overly simplistic, and clearly far from the physics fidelity required to accurately predict the performance of the plasma in a future machine, they are sufficient to design the machine coils. In addition, one can imagine that future design codes could substitute more accurate transport models into TSC in a fairly straightforward manner. For example, one could replace the Coppi-Tang transport model with the multi-scale TGLF model used in Chapter 6 of this thesis. Similarly, one could replace the simplified H-mode pedestal model in TSC with EPED [160]. In this sense, TSC resembles exactly the full tokamak simulation code that is one of the goals of modeling work throughout the fusion community, but requires many of its component models, and their interconnections, to be replaced with more accurate (and likely more computationally demanding) versions.

TSC takes as input information about the coil and vacuum vessel geometry and materials, as well as any other components in the machine which might affect the changing electromagnetic fields. It also requires coil current time traces, feedback control for coil currents based on synthetic magnetic measurements, and heating sources and their time traces, among many other things. One must also specify many other inputs, such as how to model the sawtooth crash, how to model the H-mode and the L-H transition, and others. The code outputs the plasma cross section, temperature and density profiles, fusion power and many other parameters as functions of time. The first question that one must ask is whether or not the plasma evolved stably in the machine and did not experience a vertical displacement event or some other macroscopic instability. After this, one can assess the predicted plasma performance, though, as discussed above, these predictions are far less accurate than the macroscopic stability information.

F.1 The MQ1 Tokamak

As part of this thesis work, TSC was used to design a high-field, compact, net-energy tokamak, which was named ‘MQ1,’ which stands for ‘Mission $Q > 1$.’ This tokamak was designed in a similar parameter space as the proposed SPARC device [216], which, at the time of the writing of this thesis, is still in the early design phase. Alternatively, MQ1 can be thought of as a necessary step toward the construction of a power plant like ARC [217, 218], which is based on using high temperature superconductors to increase the toroidal magnetic field in the tokamak and correspondingly reduce the machine size. The general parameters of MQ1 are also motivated by a scoping study in Reference [219].

The motivation for designing a high-field, compact tokamak like MQ1 is clear. An economically viable fusion power plant will want to maximize its fusion power

density, P_f/V_p , which is governed by the following relationship [220]:

$$\frac{P_f}{V_p} \propto \langle p \rangle^2 \propto \beta_T^2 B_0^4 \quad (\text{F.1})$$

where P_f is the total fusion power, V_p is the plasma volume, $\langle p \rangle$ is the volume-averaged plasma pressure, B_0 is the toroidal magnetic field, and $\beta_T \equiv 2\mu_0 \langle p \rangle / B_0^2$ is the toroidal plasma beta.

In this relationship, β_T is limited essentially by plasma physics (in particular MHD stability [221]), while B_0 is limited by engineering constraints. Thus, one can improve the fusion power density either by pursuing advanced plasma physics and attempting to increase β_T , or by pursuing advanced engineering solutions and increasing B_0 .

Until recently, the maximum on-axis toroidal magnetic field in a superconducting tokamak was limited to roughly 6 T by the properties of the Niobium-Tin superconducting material [219]. This limitation is one of the driving factors behind ITER's toroidal field of $B_0 = 5.3$ T [222]. In recent years, however, advances in the performance and production of Rare-Earth Barium Copper Oxide superconductors, such as Yttrium Barium Copper Oxide (YBCO), have opened up the possibility of much larger on-axis magnetic fields [216]. This new material is one of the motivations for the SPARC and ARC designs, and also for the MQ1 design presented in this appendix.

In addition to acting as an example of how one might use a full machine design code, once more accurate predictive models are developed, the design of MQ1 is used in Appendix G to evaluate the use of interferometry and polarimetry in high field machines.

Table F.1 summarizes the major machine parameters for MQ1, as well as some of the outputs of the TSC simulations. In particular, MQ1 is approximately the same size as ASDEX Upgrade, which was described in detail in Chapter 2, with a major radius of 1.65 m and a minor radius of 0.5 m, but with a magnetic field almost five times larger, 12 T on the magnetic axis. The elongation is 1.8 and the triangularity

Discharge	Value
R	1.65 m
a	0.5 m
B_0	12 T
I_p	7.5 MA
κ	1.8
δ	0.4
q_{95}	3.0
P_{aux}	30 MW
$n_{e,0}$	$4.0 \times 10^{20} m^{-3}$
$T_{e,0}$	22 keV
P_{fusion}	100 MW
Q	3.3

Table F.1: General MQ1 machine and plasma parameters, including profile and fusion power predictions from TSC.

is 0.4. The plasma current, driven entirely by the central solenoid, is approximately 7.5 MA, leading to a q_{95} of approximately 3.0. The plasma is heated with 30 MW of ICRH.

As a preview of the TSC results, the plasma is predicted to operate with a central density of $4.0 \times 10^{20} m^{-3}$ and a central electron temperature of 22 keV. This results in approximately 100 MW of fusion power and a fusion gain of 3.3.

F.2 Machine and Coil Geometry

The first input to TSC is the geometry of the machine. The toroidal field (TF) coils are not explicitly included in the simulation, since they are assumed to create a specified toroidal field that varies as $1/R$ where R is the major radius. The rest of the machine design must of course leave space for the TF coils. One must specify the locations of the central solenoid, all of the other equilibrium poloidal field coils, stability coils, a limiter, and the vacuum vessel. All of these components have an impact on what the final steady state plasma will look like, in terms of position, shape, and stability. In addition, the machine must be designed such that it is possible to get




Figure F-1: Design for the central solenoid, poloidal field coils, vacuum vessel, and limiter of MQ1. All coils are up-down symmetric and assumed to be toroidally symmetric. The toroidal field coils fit in the gap between the equilibrium poloidal field coils and the vacuum vessel.

from breakdown to the final plasma state, as this is what a real machine must do. The coils must therefore be able to shape a plasma from an initial limited configuration to a final diverted configuration. The central solenoid must also be able to drive the desired amount of current in the plasma, while maintaining the desired shape.

The central solenoid, poloidal field coils, vacuum vessel, and limiter for MQ1 are shown in Figure F-1. All of the components in this design are assumed to be toroidally symmetric. Note also that all of the components are up-down symmetric. While this configuration is not common in most present day tokamaks, a high-power compact device such as MQ1 will generate tremendous heat fluxes in the divertor. A double-null magnetic geometry partially mitigates this problem by evenly splitting the heat flux between the upper and lower divertors. While the divertor is not explicitly shown in Figure F-1, it is assumed that the divertor is toroidally continuous, which

is relevant to diagnostic design for such a machine. The limiter is only used in the ramp up portion of the discharge, before the plasma is fully developed.

The vacuum vessel is relevant to the magnetic design of the machine, as any conducting material separating magnetic coils from the plasma will cause a delay in the field propagation between the coils and plasma and vice versa. On the other hand, conducting surfaces also aid in passive plasma stability, as conducting walls resist penetration of magnetic field and thus of macroscopic plasma motion. In order to realistically simulate plasma control and stability, therefore, one must include at least an approximation of the vacuum vessel in the TSC design. In this simulation, the vacuum vessel is modeled as 1cm thick stainless steel.

The central solenoid is comprised of the OH1 and OH2 coils. OH1 is primarily used to inductively drive plasma current. OH2 is primarily used for shaping. The equilibrium field coils consist of EF1, EF2, and EF3, all used for radial force balance and plasma shaping. While these coils are called ‘equilibrium’ coils, they do have currents that vary in time, especially during the plasma ramp up in order to evolve the plasma shape, and also during the flattop in order to compensate for the current change in OH1 to drive the plasma current. The two trim coils, Trim1 and Trim2, are used for fast vertical stability control. They are located inside of the vacuum vessel in order to ensure faster field propagation to the plasma. All of the coils in the model, with the exception of the trim coils, are subdivided into multiple sections. This allows for more accurate simulation of the effect that a spatially varying magnetic field will have on the coils themselves.

Finally, the figure shows the last closed flux surface of the plasma during the discharge flattop (constant plasma current). The plasma itself is described in greater detail later, including a more detailed equilibrium in Figure F-3.

F.3 Discharge Evolution

In order to simulate the entire plasma discharge, and not just one time slice, one must specify how the coil currents evolve in time to achieve the desired discharge trajectory. One must also specify how the plasma heating is applied throughout the discharge time.

Since TSC does not simulate the breakdown event itself, the simulation starts at 0.25 s into the discharge. At this time, one specifies an initial plasma current and initial currents for each of the poloidal field coils. For the simulation in this study, the plasma begins in an inboard limited configuration with a plasma current of 250 kA. From this initial condition, one must specify the coil currents to evolve the discharge.

In particular, one inputs a desired plasma current time trace, and then assigns the task of driving the plasma current fractionally to some subset of the central solenoid and equilibrium coils. This will result in a strong swing in the coil currents in order to inductively drive the plasma current. One can also specify a PID (proportional-integral-derivative) controller to feedback control the plasma current, balancing lag between the requested and actual current, overshoot after slope changes in the time trace, and fluctuations in the current. For the simulations here, OH1 is assigned 70% of the current drive task, OH2 (upper and lower together) is assigned 22.8%, EF1 is assigned 6.3%, EF2 is assigned 0.8%, and EF3 is assigned 0.1%.

Next, one specifies time traces for the desired vertical and radial position of the plasma magnetic axis, and assigns coil groups to these tasks. The radial position of the plasma is controlled by EF3. The slow time evolution of the vertical position is controlled by EF2, and the fast feedback control of the vertical position is assigned to Trim1 and Trim2. Since the plasma in this discharge is centered at $Z=0$ (symmetric about the midplane) at all times during the simulated discharge, essentially only the trim coils are used for vertical control. This feedback is also accomplished with a PID

controller. Ensuring that a specified set of coils can achieve vertical stability is a key mission of running TSC simulations.

Finally, one controls the plasma shape throughout the discharge by specifying time traces for each of the coil sets throughout the discharge. These specified currents are added to the currents generated by the various feedback systems just described. Since these currents must be specified manually, control of the plasma shape during the course of the discharge can be rather difficult.

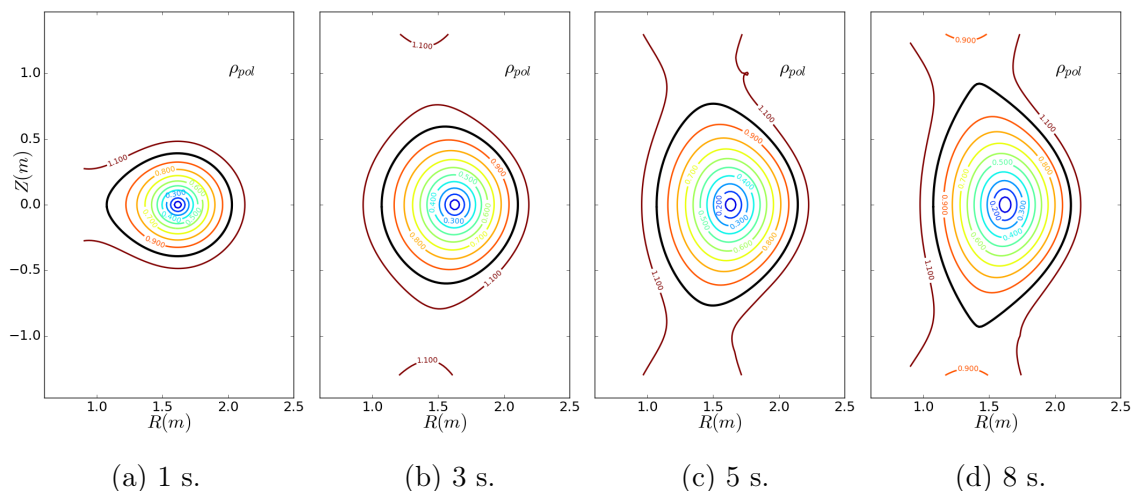


Figure F-2: Time evolution of the cross section of the plasma magnetic equilibrium. The plasma is inboard limited in the first three panels and double null diverted in the fourth panel. Contours are of ρ_{pol} , which is defined as the square root of the normalized poloidal magnetic flux. The black line is the last closed flux surface.

Overall, the MQ1 discharge is designed as follows. The plasma begins in an inboard limited configuration with a plasma current of 250 kA at 0.25 s. The plasma then approximately linearly ramps in current to achieve a current of 7.5 MA at 7.75 s. During this time the plasma is slowly enlarged in cross-sectional area and shifted from an inboard limited to a double null diverted configuration. This process is illustrated in Figure F-2. There is then a current flattop of 12.25 s, until discharge time 20 s, which is discussed further in the next section about peak plasma performance. The plasma is then ramped down in current before ending the discharge.

During this time the plasma heating also evolves. The discharge begins with no

auxiliary heating (it is heated ohmically). At 1s, the 2.5 MW of ICRH heating power is applied. At 3 s, this is increased to 5 MW, and then to 7.5 MW at 5s. At 7.75 s, at the end of the current ramp, the heating power is increased to 30 MW. This final increase in heating power is also assumed to trigger the transition of the plasma to H-mode, inducing an edge pedestal in the temperature and density profiles.

F.4 Predicted Peak Plasma Performance

After the current ramp, the plasma reaches peak performance and remains relatively steady throughout the flattop portion of the discharge. This is the portion of the discharge that is of interest for fusion power production.




Figure F-3: A cross section of the magnetic equilibrium of MQ1 during flattop operation. Contours are of ρ_{pol} . The black line is the last closed flux surface.

In terms of plasma shape and geometry, Figure F-3 shows the magnetic equilibrium during the flattop portion of the discharge at 12 s. This is very similar to the

equilibrium from Figure F-2 (d) at 8 s, differing only due to a slight relaxation of the plasma to an equilibrium current profile. In this flattop configuration, the plasma has a total current of 7.5 MA and an edge safety factor q_{95} of 3.0. The safety factor is below 1.0 on the magnetic axis, so the plasma is sawtoothing. With the caveat that the TSC sawtooth model is only very basic, the code predicts a sawtooth period of approximately 2.5 seconds and an inversion radius of approximately $\rho_{pol} = 0.15$.

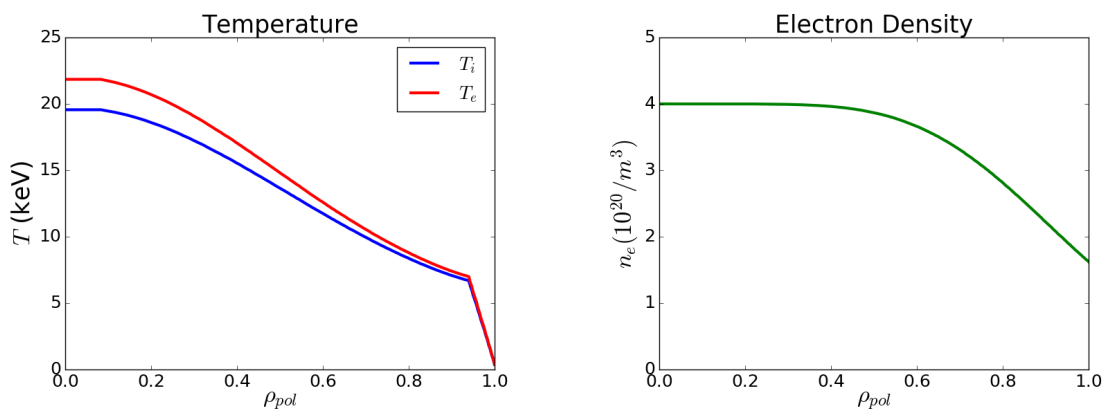


Figure F-4: Predicted ion (blue) and electron (red) temperature (a) and electron density (b) profiles for MQ1 during flattop operation.

During the flattop portion of the discharge, TSC predicts the temperature and density profiles shown in Figure F-4. As stated above, these profiles are predicted using the Coppi-Tang transport model. The model includes the effect of alpha particle heating from fusion, and assume a 50-50 mix of deuterium and tritium to calculate the total fusion power. The plasma effective charge Z_{eff} is calculated to be approximately 1.2, once one includes the presence of helium ash and other impurities.

These profiles predict a total fusion power of approximately 100 MW. Given the ICRH heating power of 30 MW (and the fact that ohmic heating is nearly negligible, less than 1MW, during the flattop), this means that MQ1's predicted fusion gain is $Q \approx 3.3$.

Conclusions

This appendix described TSC, a fully integrated simulation code used primarily to design the poloidal field coils for tokamaks. While many aspects of tokamak physics have been considerably simplified in TSC, the code serve as an example of what an integrated tokamak simulation code might eventually look like. By replacing the turbulent transport model in TSC with TGLF (or even a full gyrokinetic code), and replacing each of the other models used in TSC with higher fidelity versions, one would eventually be able to fully simulate an entire tokamak discharge, from breakdown to ramp down, and predict the performance of the machine before building it. Validation, which has been the main focus of this thesis, ensures that each of the models used in a code like TSC are simulating the plasma accurately.

Appendix G

InterPol Interferometer Polarimeter

This appendix presents a conceptual design for a combined interferometer and polarimeter, called InterPol, which will measure both line-integrated density and poloidal magnetic field in a compact, high-field, high-density, net energy tokamak. InterPol is designed to work on the MQ1 tokamak, modeled with TSC in Appendix F. The high field and high density of MQ1 offer unique opportunities for interferometry and polarimetry diagnostics, allowing one to measure both density and poloidal magnetic field (which is related to the plasma current profile) with a single set of CO_2 and $HeNe$ lasers. Unlike other diagnostics described in this thesis, InterPol is intended primarily as an equilibrium diagnostic, measuring the steady state and slow time evolution of the density and poloidal magnetic field, though this appendix will show that it will likely be able to resolve some fluctuations as well, for example from toroidal Alfvén Eigenmodes. The remainder of this appendix introduces interferometry and polarimetry in greater detail (Section G.1), outlines the InterPol design (Section G.2), and presents predicted signals from synthetic diagnostic calculations on the TSC outputs (Section G.3). The work of this appendix is pending publication as Reference [223].

G.1 Interferometry and Polarimetry

This thesis has already discussed a few diagnostics in detail, such as electron cyclotron emission and reflectometry, and has mentioned many others, such as Thomson Scattering, CXRS, etc. This appendix is concerned with two other diagnostics, interferometry and polarimetry.

Interferometry is used to measure line-integrated electron density in a tokamak, and polarimetry is used to measure the line-integrated product of density and magnetic field parallel to the beam direction [66]. This section describes the governing principles of these two diagnostics in order to then explain why the diagnostics are different on a high-field, high-density tokamak and to make predictions of signal levels.

The physical principle behind interferometry is based on the phase delay of a beam of light traveling through a plasma compared to a beam of light traveling through a vacuum. To first order in electron temperature, the phase delay is given by [66, 224]:

$$\Delta\phi = \frac{-e^2\lambda}{4\pi c^2 m_e \epsilon_0} \int n_e \left[1 - \frac{3}{2} \frac{T_e}{m_e c^2} \right] dl \quad (\text{G.1})$$

where $\Delta\phi$ is the phase delay in the beam, e is the fundamental charge, λ is the wavelength of the beam, c is the speed of light, m_e is the electron mass, ϵ_0 is the permittivity of free space, n_e is the local electron density, T_e is the local electron temperature, and dl is the path length of the beam.

In present day tokamaks, the temperature term in Equation G.1 is often negligible, in which case the phase delay can be used directly to calculate the line-integrated density. Generally, in order to account for the temperature term, one would have to have independent knowledge of the electron temperature profile. As will be discussed later, however, it is possible to use data from a combined interferometer-polarimeter, such as the diagnostic discussed in this appendix, to account for finite temperature effects without input from a separate diagnostic.

Interferometry has a long history of tokamak measurements, and has become an incredibly robust diagnostic. It is able to measure line-integrated density quickly and without complicated analysis, allowing for active feedback control of the plasma density [35]. In addition, one can use multiple chord measurements in order to invert an approximate density profile. Interferometer measurements are vital to machine operation and optimization of plasma performance.

Polarimetry operates on a similar physical principle, utilizing the rotation in the polarization of a linearly polarized wave as it propagates through a plasma parallel to a magnetic field. This effect is known as Faraday Rotation [66]. Including again a finite temperature correction, the Faraday Rotation angle is given as [66, 224]:

$$\alpha = \frac{e^3 \lambda^2}{8\pi^2 m_e^2 c^3 \epsilon_0} \int n_e \left[1 - 2 \frac{T_e}{m_e c^2} \right] \vec{B} \cdot d\vec{l} \quad (\text{G.2})$$

where α is the Faraday Rotation angle (half of the angle by which the linear polarization rotates), \vec{B} is the background magnetic field, $d\vec{l}$ is the beam path, and all other quantities are defined as above.

As for interferometry, the temperature term is often negligible in present day tokamaks, but will not be in a net-energy machine such as MQ1. One can use the results of interferometry to calculate the plasma density, and thus relate the Faraday Rotation to the line-integral of the magnetic field parallel to the beam propagation. Specifically, a radially viewing polarimeter measures the poloidal magnetic field, which can be related to the safety factor and current profiles. Knowledge and manipulation of these profiles is important to both plasma stability and optimization of future steady state machines [225].

In addition to slowly evolving equilibrium profiles, both diagnostics are also able to measure plasma fluctuations (in density and magnetic field respectively). These measurements may be useful for turbulent transport validation, like the CECE measurements from Chapter 4, as well as studying other plasma fluctuations, such as

Alfvén Eigenmodes [226, 227].

As mentioned above, while the temperature terms in Equations G.1 and G.2 are typically negligible in present day tokamaks, the significantly higher temperatures in any net-energy machine, such as MQ1 or ITER, mean that these terms will be important. Fortunately, by using a combined interferometer and polarimeter, it is possible to correct for finite temperature effects to first order without using data from a separate electron temperature diagnostic. Details of these calculations can be found in References [224, 228]. This technique does not correct second order effects, but these are typically of 0.1% or less, which is significantly less than the instrumental uncertainty for polarimetry measurements, and is similar to the instrumental uncertainty for interferometry measurements. If desired, these second order effects can be corrected using independent temperature profile measurements.

Since part of the motivation for designing InterPol on MQ1 was the higher magnetic field and correspondingly higher plasma current and density, consider the effect that this has on interferometry and polarimetry. As shown in Equation G.1, the signal level in interferometry is proportional to density, and as shown in Equation G.2, the signal level in polarimetry is proportional to both density and poloidal magnetic field (which depends on plasma current density). For this reason, a high-field, high-density device like MQ1 should expect larger signal levels for both diagnostics than a similarly sized low-field device. This enables the use of shorter wavelength beams for both diagnostics, for which more reliable sources exist. In particular, InterPol will use a single set of CO_2 and $HeNe$ for both interferometry (two color interferometry) and polarimetry. This design choice is discussed further in Section G.2.

On the other hand, a compact size also leads to challenges for these diagnostics. In particular, producing a large amount of fusion power in a compact device leads to large neutron fluxes, requiring more careful shielding of neutron-sensitive components. In addition, as discussed in Appendix F, high divertor heat fluxes will likely require a

double null configuration with toroidally continuous divertors. This precludes vertical lines of sight that are common on other machines.

Using the fundamental principles behind interferometry and polarimetry described in this section, a conceptual design for InterPol is now described.

G.2 Proposed InterPol Diagnostic

This section outlines the design of InterPol on MQ1, using the equilibrium from Appendix F and addressing the challenges and opportunities associated with a high-field compact device discussed in Section G.1.

Starting with the lines of sight for InterPol, many existing tokamaks, including Alcator C-Mod [35] and EAST [229], use a series of parallel beams. Using a series of parallel beams typically simplifies engineering of the diagnostic, and also makes the inversion of line-integrated data to obtain radial profiles more straightforward. In a device running pure deuterium plasmas, this geometry is acceptable, but in a device producing a considerable neutron flux, one would like to minimize penetrations in the first wall. A series of parallel beams requires many penetrations in the first wall of the machine, but it is possible to instead use a single penetration and arrange the beams in a fan.

It is for this reason that InterPol proposes using a single penetration at the low-field side midplane for all of the beam chords, similar to the poloidal polarimeter design on ITER [230]. Figure G-1 shows this chord geometry overlaid on the MQ1 plasma density cross section from TSC. Density is given in units of 10^{20} m^{-3} , and the red lines are the proposed beam chords.

As shown in Figure G-1, the system consists of 11 chords, originating from a single penetration in the outboard midplane and propagating inward at angles of 0° , 7° , 14° , 21° , $\pm 27^\circ$, 32° , 37° , 41° , and $\pm 45^\circ$. Each beam reflects off of a retroreflector on the high-field side of the machine, before passing through the plasma again and exiting




Figure G-1: InterPol chord geometry overlaid on a density cross section of MQ1. Red lines indicate the InterPol laser chords. Density contours are in units of 10^{20} m^{-3} . An asymmetric chord arrangement gives greater maximum spatial resolution without an excessive number of chords, and works best if the plasma is up-down symmetric.

through the same penetration in the first wall. These retroreflectors will be sunk in to the first wall in order to minimize direct neutron exposure. The chord angles were chosen to cover as much of the plasma volume as possible while still using only a single first wall penetration and without requiring retroreflectors in the divertor region. The up-down asymmetric arrangement was chosen to minimize the number of interferometer chords, while still offering high radial resolution in some regions of the plasma.

After selecting the diagnostic geometry, one must choose the laser wavelength for the measurements, carefully balancing several competing effects. Increasing the laser wavelength increases signal levels for both interferometry and polarimetry. Equation

G.1 reveals that the interferometer phase shift is proportional to λ , and Equation G.2 shows that the polarimeter Faraday Rotation is proportional to λ^2 . In addition, using longer wavelengths decreases the effect of mechanical vibrations, which are more detrimental at short wavelengths since the resulting noise depends on the ratio of the path length change to the laser wavelength. On the other hand, the Cotton-Mouton effect, which results in an initially linearly polarized wave becoming elliptically polarized as it propagates through the plasma and complicates polarimetry measurements, depends on λ^3 [66], and so longer wavelengths are more problematic. In addition, refraction of the laser beam through the plasma also increases with increasing wavelength, which can be problematic. Finally, one must consider the availability of reliable laser sources at a given wavelength, since reliability is one of the reasons that interferometry in particular is key to tokamak operation.

Taking all of these considerations into account, InterPol proposes using 10.6 μm wavelength CO_2 lasers for the primary interferometry and polarimetry measurements and 0.633 μm $HeNe$ for vibration correction, in what is known as two-color interferometry (TCI) [35]. The phase change in the CO_2 laser is primarily sensitive to the plasma, but will also have some contribution from vibration. On the other hand, due to the very short wavelength of the $HeNe$ laser, the change in phase is completely dominated by mechanical vibration, with the plasma contribution almost negligible. For this reason, one can use the $HeNe$ signal to correct the CO_2 phase delay for mechanical vibration, resulting in a measurement purely of the plasma density. Mathematically, one can write [35]:

$$\int n_e \left[1 - \frac{3}{2} \frac{T_e}{m_e c^2} \right] dl = \frac{4\pi c^2 m_e \epsilon_0}{e^2} \frac{\Delta\phi_{HeNe} \lambda_{HeNe} - \Delta\phi_{CO_2} \lambda_{CO_2}}{\lambda_{CO_2}^2 - \lambda_{HeNe}^2} \quad (\text{G.3})$$

This is the same combination of lasers used for TCI on Alcator C-Mod [35]. Most machines, however, have considerably lower density than Alcator C-Mod, and so must use longer wavelengths to achieve reasonable signal levels, such as $\lambda = 119 \mu\text{m}$ far

infrared (FIR) CH_3OH lasers on the Large Helical Device [210] (also described in Appendix E) and $\lambda = 195 \mu\text{m}$ deuterium cyanide (DCN) lasers on ASDEX Upgrade [44, 231] and JET [232]. Section G.3 will show that the signal level from CO_2 lasers will be sufficient in the high density in MQ1. One of the main advantages of using CO_2 lasers is that they are incredibly reliable compared to other sources. For example, the CO_2 lasers in the ITER tangential interferometer-polarimeter (which is discussed in further detail below) are expected to fail in less than one discharge per year [233].

The consideration of reliability is one of the primary reasons that CO_2 lasers are preferred to an intermediate wavelength system with $\lambda \approx 50 \mu\text{m}$, which might allow more accurate measurements while keeping the Cotton-Mouton effect and refraction at acceptable levels (these are discussed below for the CO_2 lasers). While there has been some development of a two-color interferometer-polarimeter based on $57.2 \mu\text{m}$ and $47.7 \mu\text{m}$ CH_3OD sources [234, 235], this technology is still immature compared to CO_2 lasers.

Turning now to polarimetry, the choice of wavelength is more constrained than for interferometry due to the λ^2 dependence of the Faraday Rotation angle. This dependence means that the signal levels from short wavelength lasers are generally far too small on present day machines, and that most current system use wavelengths longer than $100 \mu\text{m}$ for poloidal polarimetry. For example, the same $\lambda = 195 \mu\text{m}$ DCN lasers are used for both interferometry and polarimetry on ASDEX Upgrade [44, 231], and the combined polarimeter-interferometers on both EAST (POINT [229]) and J-TEXT (POLARIS [236]) use $\lambda = 432.5 \mu\text{m}$ formic acid ($HCOOH$) FIR lasers.

All of the discussion of polarimeter wavelengths up to this point has been concerned with poloidally viewing geometries, which measure the poloidal magnetic field. Tangentially viewing polarimeters, on the other hand, have different constraints. Since polarimetry is sensitive to the magnetic field parallel to the beam propagation direction, tangentially viewing polarimeters are sensitive to toroidal field, which

is generally much stronger than the poloidal field. This, often combined with a longer path length, means that tangential polarimeters can use shorter wavelengths even in low density and low field devices, such as the 10.6 μm wavelength CO_2 lasers on the DIII-D interferometer-polarimeter [237]. Similarly, the proposed tangential interferometer-polarimeter (TIP) on ITER will also use CO_2 lasers [228, 238]. These systems have a fundamentally different purpose than poloidal polarimeters, however, since they do not measure the poloidal magnetic field. The Faraday Rotation measurement is instead used as a second, independent measure of the line-averaged density, since the toroidal magnetic field is generally well known.

In contrast, MQ1's high density and high poloidal magnetic field allow for the novel use of 10.6 μm CO_2 lasers for both interferometry and poloidal polarimetry. Section G.3 calculates the anticipated signal levels using a synthetic diagnostic, showing them to be clearly measurable.

Finally, while polarimetry cannot account for the effects of mechanical vibration in the same way that two-color interferometry does, various methods have been developed to minimize these effects. Most simply, one can minimize the vibrations themselves by carefully designing and placing optical components in a vibration damped environment [229]. One can also use more clever detection techniques, such as using the signal from a detector that measures the beam phase before it passes through the plasma as the local oscillator in the final detector, as opposed to using a separate local oscillator. This has been shown to considerably reduce the vibration-induced error in polarimetry [228, 238].

G.3 Predicted Signals and Measurement Capabilities

In order to determine whether or not the proposed InterPol design will be able to make useful measurements on MQ1, one must estimate the expected signal levels

Chord Angle (Degrees)	Interferometer Phase Shift $\Delta\phi$ (Degrees)	Polarimeter Faraday Rotation α (Degrees)
Center (0)	1166	0.00
7	1162	0.39
14	1150	0.75
21	1128	1.06
27	1091	1.29
32	1038	1.42
37	972	1.46
41	900	1.43
45	824	1.36

Table G.1: InterPol signal levels during flattop operation of MQ1, as calculated using the proposed lines of sight and the 2D MQ1 density and poloidal magnetic field data. Two pass signal levels are shown for all of the upper chords of InterPol.

using a synthetic diagnostic and the MQ1 profiles from Appendix F. In particular, this analysis uses the 2D density profiles and diagnostic lines of sight from Figure G-1, as well as the 2D poloidal magnetic field cross section from the MQ1 magnetic equilibrium.

Equilibrium Measurements with MQ1

The predicted interferometry and polarimetry signal levels are evaluated as follows. For interferometry, the phase shift was calculated by line-integrating the density along the chosen lines of sight in the 2D MQ1 geometry. This line-integrated density was then inserted into equation G.1 in order to get the total phase shift. Similarly for polarimetry, the predicted Faraday Rotation was calculated by integrating the density times the poloidal magnetic field dotted with the beam path, and then inserting this into Equation G.2. The predicted signal levels calculated in this manner for two-pass operation (after the beam has passed through the plasma, reflected off of a retroreflector, and passed back through the plasma) during the flattop portion of the

MQ1 discharge (at 12 seconds into the discharge) are shown in Table G.1 for all of the upper chords.

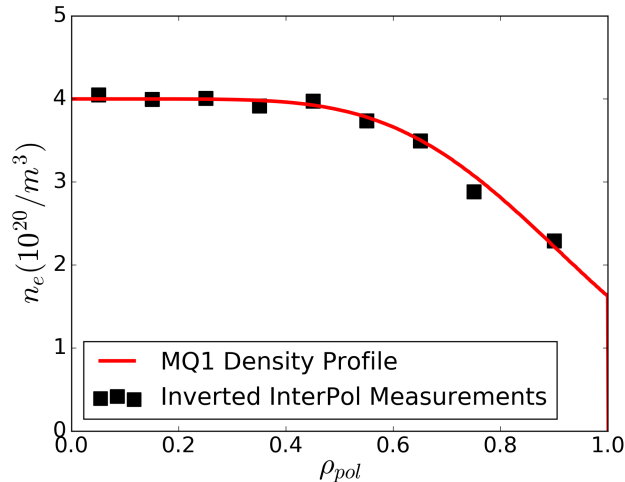


Figure G-2: MQ1 density profile and synthetic measurement with InterPol. The InterPol profile was calculated from the upper chord line integrated measurements using standard inversion techniques.

To put these signal levels into context, existing systems have been able to obtain instrumental resolution as small as $\Delta(\Delta\phi) = 1.5^\circ$ for the interferometer phase shift and as small as $\Delta\alpha = 0.06^\circ$ for the polarimetry Faraday Rotation [228]. For the MQ1 flattop signal levels shown in Table G.1, this resolution corresponds to an accuracy of between 0.1% and 0.2% for the density measurement and between 4% and 8% for the poloidal magnetic field measurement. This suggests that the diagnostic will be able to make high quality measurements of both density and poloidal magnetic field during peak plasma performance. To illustrate this point, Figure G-2 shows the MQ1 density profile from TSC, as well as a profile that has been inverted from the synthetic InterPol measurements. As this figure shows, InterPol’s chord geometry enables high resolution density profile measurements using standard inversion techniques.

While the signal levels shown in Table G.1 were for the flattop portion of the MQ1 discharge, earlier in the discharge both the density and poloidal field will be lower, and the plasma will be smaller in cross-sectional area (as shown in Figure F-2). For

this reason the expected signal levels will be smaller, and the relative errors will be larger. Very early in the current ramp (at 1 s into the discharge), the relative error on some chords may be as much as 10% for interferometry and 100% for polarimetry.

Finally, three possible issues are briefly addressed. First, the issue of fringe jumps in the interferometer signal, which may complicate interpretation of the data at large phase angles [66], can be resolved by using the polarimeter measurement as a reference. Since the Faraday Rotation angle will always be much greater than π (as shown in Table G.1), the signal can be used to determine the direction of change in the interferometer phase angle when fringe jumps occur near multiples of π .

Second, the short wavelength of the CO_2 lasers in InterPol mean that the Cotton-Mouton effect will be negligible in this system. Specifically, the ellipticity introduced by the Cotton-Mouton effect will be less than 10^{-3} .

Finally, the short wavelength also means that refraction will be negligible. Ray tracing calculations reveal that refraction will cause the beam path to deviate from a straight line by less than 0.1° , which is comparable to or less than the feasible beam alignment accuracy and the effects of vibrations.

InterPol Fluctuation Measurements

While most of this appendix has been devoted to discussing the use of InterPol as an equilibrium diagnostic, functioning to feedback control the plasma density and measure the current profile, InterPol will also function as a fluctuation diagnostic. Fluctuation measurements depend quite strongly on the precision and spatial and temporal resolution of the diagnostic, and so some additional attention is merited.

One of the major concerns for any fluctuation diagnostic is the noise level. These noise levels are typically different than the instrumental resolution that was discussed earlier in this section. In particular, below 5 kHz, vibration and other effects often limit the ability to perform fluctuation measurements. For $f > 20$ kHz, however,

an existing CO_2 interferometer-polarimeter found that the noise levels for both interferometry and polarimetry were approximately $[10^{-6} - 10^{-5}](^\circ)^2\text{kHz}^{-1}$ [238]. This would enable fluctuation measurements down to approximately $\tilde{n}/\bar{n} \sim 10^{-5} - 10^{-6}$ for density and $\tilde{B}_p/\bar{B}_p \sim 10^{-2} - 10^{-3}$ for magnetic field on MQ1 (where the bars represent a line average and the tildes represent the fluctuating quantity). Based on what has been observed on current machines, these noise levels should allow for precise measurements of density fluctuations from MHD, turbulence, and other sources, and marginal measurement of magnetic field fluctuations from similar sources [238].

In addition to noise levels, temporal resolution is important to the capability of fluctuation diagnostics, as some fluctuations in tokamaks exist at quite high frequencies. Existing interferometer-polarimeter systems have achieved time resolutions of up to 500 kHz [238] and a pure polarimeter system has operated at up to 1 MHz [226]. This level of time resolution has been sufficient to measure a wide range of tokamak phenomena, including Alfvén Eigenmodes [228, 238], the quasi-coherent modes (QCMs) typical of enhanced D_α (EDA) H-modes on Alcator C-Mod [226, 227, 239], and broadband turbulent fluctuations [239].

To choose just one example of a fluctuating mode that may be of interest on MQ1, one can estimate the frequency of toroidal Alfvén Eigenmodes (TAEs) for the MQ1 parameters. The typical TAE frequency is given by [240]:

$$f_A \approx \frac{v_A}{4\pi q R} \approx \frac{B_0}{4\pi q R \sqrt{\mu_0 \sum n_i m_i}} \quad (\text{G.4})$$

where v_A is the Alfvén speed, q is the local safety factor, R is major radius, and $\sum n_i m_i$ is the mass density of the plasma. For MQ1, one estimates a TAE frequency of approximately 650 kHz. Assuming that one can digitize at a rate of 1 MHz as in Reference [226], it should be possible to resolve this mode. In addition, it may also be possible to tailor specific discharges to have lower frequency modes that would be more easily measurable.

Finally, spatial resolution is important for fluctuation diagnostics, especially those concerned with turbulence measurements. In particular, for fluctuation measurements one is concerned with both the width of each beam and the spacing between beams.

First, the beam diameter w_0 of the laser in the plasma determines the minimum measurable fluctuation wavelength, as any fluctuation smaller than the beam width will be averaged out in the measurement. For this reason, the maximum measurable wavenumber for a given beam width is $k_{max} = 2/w_0$ [238]. An existing diagnostic using CO_2 lasers on DIII-D has achieved a beam radius of approximately 9 mm, allowing measurements up to wavenumbers of approximately 1 cm^{-1} [228]. A similar system on Alcator C-Mod has been able to measure wavenumbers up to 3.2 cm^{-1} [239]. With this maximum resolvable wavenumber it is possible to conduct comprehensive studies of both density and magnetic field fluctuations from a variety of sources [238, 239].

The spacing between beams is also important for fluctuation measurements, as it sets the spatial wavenumber resolution. In particular, a beam separation of ΔL gives a wavenumber resolution of $\Delta k = 2\pi/N\Delta L$ where N is the total number of chords [239]. Thus a larger number of chords spaced more closely together gives a finer wavenumber resolution, motivating the asymmetric chord geometry in Figure G-1. On the other hand, it is certainly possible to make non-wavenumber resolved fluctuation measurements, which has been the case in most past studies [226, 227, 238].

These spatial constraints are typically much more stringent than for equilibrium measurements, as density feedback control really only requires one chord, and profile reconstruction requires more chords, but places few restrictions on beam width.

Conclusion

This appendix has presented the proposed InterPol diagnostic, a combined interferometer-polarimeter for the MQ1 tokamak described in Appendix F. This diagnostic takes

advantage of the high density and high poloidal magnetic field on MQ1 in order to perform both interferometry and poloidal polarimetry measurements with a single set of CO_2 and $HeNe$ lasers. The design also conforms to the constraints of a high neutron and heat flux environment, launching all of its chords through a single first wall penetration at the outboard midplane and avoiding any impediment to toroidally continuous upper and lower divertors. These measurements will be important for density feedback control on MQ1, as well as enabling density and poloidal magnetic field profile measurements. In addition to equilibrium measurements, InterPol will also be able to measure fluctuations due to turbulence and Alfvén Eigenmodes, among other things. Synthetic diagnostic testing using MQ1's profiles suggests that InterPol will be capable of all of these measurements, in terms of signal level, noise level, temporal resolution, and spatial resolution. These measurements would both ensure successful machine operation and contribute another constraint for future turbulent transport model validation studies such as those performed in this thesis.

Appendix H

Specifics of Data Analysis and Simulation Workflows

This appendix contains information about the specific workflows and codes used for some of the data analysis in this thesis. It is primarily of interest to the reader that will be working with Alcator C-Mod and ASDEX Upgrade data. Some references are made to particular figures in the main body of the thesis, describing how these figures are made.

H.1 Perturbative Diffusivity Analysis

This section describes how to perform the partial sawtooth heat pulse analysis to measure the perturbative thermal diffusivity on Alcator C-Mod and ASDEX Upgrade. The theory and results of this analysis are shown in Chapter 3 and are used in Chapters 5 and 6.

1. Identify Partial Sawtooth Crash Times

Identification of partial sawtooth crashes must be performed manually on both Alcator C-Mod and ASDEX Upgrade at the moment. Automating this process would be a potential future project related to this work.

Identification of partial sawtooth crash times is done using time series electron temperature data. On Alcator C-Mod, access GPC data stored in MDSplus, using the ‘scopes’ code, or whichever data reader the user prefers. On ASDEX Upgrade, access the 1D ECE radiometer (the diagnostic name is ‘CEC’ in the shotfile system) using ‘cview’ or whichever data reader the user prefers. Load the desired discharge’s data.

As described in Chapter 3, partial sawteeth are characterized by a drop in temperature slightly off axis in an annular region in the tokamak, without a drop in core temperature. One should therefore look through the time traces to find a time a time where the temperature drops sharply on a slightly off-axis channel, but does not change on axis. Record this time for use in the code. Go through the entire discharge (or time period of interest), identifying all of the partial sawtooth crashes.

2. Insert Times into Perturbative Diffusivity Analysis Code

Open the analysis code for the tokamak of interest.

For Alcator C-Mod, open the MATLAB script ‘/home/creely/matlab/ AnalyzeShotAux2.m’ in the MIT PSFC server. Enter the desired discharge number. Enter all of the partial sawtooth times in either the ‘tStartL’ or ‘tStartI’ arrays. Ignore the other array if only one time period is of interest. Enter the GPC channel of the inner limit of the analysis as ‘inChannel’ and the outer limit channel as ‘outChannel.’ The innermost channel should be first GPC channel outside of the sawtooth mixing radius, and the outermost channel should be the last channel on which the heat pulse is still distinguishable from noise. These are typically channels 5 and 7 or channels 6 and 8, depending on the location of the mixing radius.

For ASDEX Upgrade, use the Python script ‘/afs/ipp-garching.mpg.de/home/a/

acreely/python/shotChiPertpy6.py' from the command line, using the following syntax: `python shotChiPertpy6.py shot 'pulseTimes' inChannel outChannel timeRange smooth backgroundAvg`. Pulse times are entered as space separated values in a string. The 'inChannel' and 'outChannel' inputs may have to be adjusted based on the discharge, as the ECE radiometer channel locations are occasionally changed on ASDEX Upgrade. The numbers represent the order from the center of the plasma, not the recorded channel number, as the recorded channel numbers are not necessarily in order in the plasma. A typical channel range is 11 to 14.

3. Adjust Code Parameters if Necessary

Run the code. This will output the average perturbative diffusivity, an uncertainty, the perturbative diffusivity for each individual heat pulse, and a number of other plasma parameters, such as density, current, etc. These are not yet the final results. This will plot all of the channels for each sawtooth pulse in a different figure. For Alcator C-Mod, adjust the 'LRange' or 'IRange' parameters, which set the length of the time window that is analyzed around each pulse, such that the window is long enough to capture the peak of the heat pulse on the outermost channel but short enough that the time window does not include the next sawtooth crash. Adjust 'LSmooth' and 'ISmooth' to sufficiently smooth noise but to still resolve the fast rise in the innermost channel.

For ASDEX Upgrade, adjust the 'timeRange' and 'smooth' parameters in the same way as the 'LRange' and 'LSmooth' parameters on Alcator C-Mod.

4. Check Analysis for Each Heat Pulse

Rerun the code once parameters have been adjusted. This is still not the final result. Go through each heat pulse figure and check to make sure that the data is clean enough for inclusion in the analysis. This process is essentially the

same on both machines. You may have to adjust slightly the crash time for each individual pulse so that the window includes the peak on the outermost channel but excludes the next heat pulse. Remove heat pulses (by deleting their time from the list in the code) that have noise levels significantly higher than others, or which appear to have double sawtooth crashes or other odd events that are not clean partial sawteeth. Also remove pulses for which the smoothing algorithm has over- or under-smoothed the data. It is also possible, at the user's discretion, to remove pulses that have a perturbative diffusivity which is a significant outlier (as shown in the code output).

5. Run Analysis Code with Final Parameters and Pulse Times

Once the analysis parameters have been adjusted and bad heat pulses have been removed, run the code a final time to obtain the final perturbative diffusivity and uncertainty. These will print in the command line. Note that the ASDEX Upgrade code will also output the perturbative diffusivity calculated for a composite heat pulse, as described in Chapter 3. These final code results were used to generate, for example, Figure 3-3.

6. Save Output if Desired

If desired, save the results for larger data set studies. On Alcator C-Mod, just save the MATLAB workspace. On ASDEX Upgrade, call the `chiCalc.writeOutputs()` function within the Python script to save the outputs to a text file.

H.2 CECE Data Analysis

This section describes the analysis routines used for CECE data. The theoretical basis for this analysis is given in Chapter 4 and Appendix C. These data analysis procedures were used for all of the results in Chapters 4, 5, and 6.

1. Select Analysis Time Range

The first step in analyzing CECE data is to select an appropriate discharge with a long steady phase of plasma operation, as longer time averaging significantly reduces the uncertainty and sensitivity limit. Typically one should aim for an absolute minimum of 150 ms of steady operation, and preferably up to a few seconds. One must also ensure that the toroidal magnetic field is chosen such that the ECE resonance location of the CECE frequencies will be at the desired location in the plasma. This can be estimated analytically, and is calculated more precisely within the analysis code.

2. Run CECE Analysis Code to See Spectra

Once the raw data has been collected, one runs the data analysis code a first time to see the fluctuation spectra and determine the frequencies over which to integrate the turbulent feature and the frequency range to use for the background subtraction.

On Alcator C-Mod, use the Python script `‘/home/creely/python/fftcece15cmod.py’` from the command line with the following syntax: `python fftcece15cmod.py shot ‘channels’ time1 time2 fftBin lowpassFreq fluctFreqLow fluctFreqHigh backFreqLow backFreqHigh`. In this call, ‘shot’ is the plasma discharge for analysis, ‘channels’ is a space separated string of the channels to be analyzed (typically ‘01 02 03 04’ or ‘01 02 03 04 15 16’ depending on which system was installed when the discharge was taken), ‘time1’ and ‘time2’ are the start and end times for the analysis, ‘fftbin’ is the number of points to use in the ensemble average Fourier transform (typically 512 or 1024), ‘lowpassFreq’ is the frequency of the digital low pass filter in MHz (typically 1MHz), ‘fluctFreqLow’ and ‘fluctFreqHigh’ set the lower and upper frequency bounds (in kHz) for integration of the total temperature fluctuation level, and ‘backFreqLow’ and ‘backFreqHigh’ set the lower and upper frequency bounds (in kHz) for the background level

assessment.

The command line call on ASDEX Upgrade is identical, using the code located at ‘/afs/ipp-garching.mpg.de/home/a/acreely/jupyter/fftcece15aug.py.’ One can also use the interactive Jupyter Notebook version, for example at ‘/afs/ipp-garching.mpg.de/home/a/acreely/jupyter/Multi_Machine_CECE.ipynb.’

This initial call of the code will generate fluctuation spectra as well as radial profiles, but at this point only the spectra are of interest.

3. Determine Frequency Range for Fluctuation Integration and Background Subtraction

Once the initial spectra have been generated, one can determine the frequency ranges for the total fluctuation level integration and for the background subtraction. At this point, one should also identify whether there are any coherent features in the spectrum that are non-turbulent phenomena, such as MHD modes or other features. These can be notch filtered out if necessary.

Look at the spectrum, and if there is a clear turbulent feature, note the lower and upper frequency bounds of the feature. These will be ‘fluctFreqLow’ and ‘fluctFreqHigh’ in the next run of the analysis code. Next, identify a frequency range that is clearly above the turbulent feature, such that the complex coherence function is essentially flat with frequency, except for the contributions from noise. The bounds of this region will be used as ‘backFreqLow’ and ‘backFreqHigh’ in the next run of the code.

4. Run CECE Analysis Code for Final Results

Once the frequency ranges of the turbulent feature and the background have been set, rerun the code for the final analysis results. This run of the code should produce spectra and a radial profile, as well as printing the total temperature fluctuation level, uncertainty, sensitivity level, and radial location for

each channel pair. The code itself can be modified to plot different combinations of spectra.

The Alcator C-Mod code was used to generate, for example, Figures 4-8 and 4-9, and the ASDEX Upgrade code was used to generate Figure 4-11.

H.3 TRANSP

There are numerous descriptions on how to run TRANSP, so this section will not attempt to describe this whole process. Instead, this section focuses on how to run a few post-processing routines that were written as part of and used in this thesis.

On Alcator C-Mod, the code is located at `‘/home/creely/python/transpRead.py’` and on ASDEX Upgrade it is located at `‘/afs/ipp-garching.mpg.de/home/a/acreely/python/transpReadAUG2.py’`.

1. Run TRANSP and Download Data from TRANSP Server

After running TRANSP on the servers at PPPL, one must download the data to the local server at either Alcator C-Mod or ASDEX Upgrade. At Alcator C-Mod, this is done automatically, and one can then access the data through the MDSplus data storage system. At ASDEX Upgrade, this must be done manually. One must fetch the data from the PPPL servers, unpack it, and then write a shotfile to store it in the user’s personal directory. For this thesis, this was done using the TRGUI program. Note that on ASDEX Upgrade TRANSP runs are not publicly available to other users, and that one must either repeat a run or be granted access to another user’s directory in order to access the TRANSP data.

2. Run TRANSP Plotting Script

Once the data is either stored in MDSplus (at Alcator C-Mod) or as a shotfile

Number	String Input	Quantity to Plot
1	profiles	Temperature and Density Profiles
2	gradients	Normalized Gradient Scale Lengths
3	q	q Profile
4	Er	Radial Electric Field
5	diffusivities	Power Balance Thermal Diffusivities
6	powerbale	Volumetric Electron Power Balance
7	powerbali	Volumetric Ion Power Balance
8	powerbaletot	Total Electron Power Balance
9	powerbalitot	Total Ion Power Balance
10	heatflux	Electron and Ion Heat Fluxes
11	heatfluxe	Electron Heat Flux
12	heatfluxi	Ion Heat Flux
13	neutrons	Neutron Count (against time)
14	energy	TRANSP Stored Energy (against time)
15	current	Plasma Current (against time)
16	voltage	loop voltage
17	energyMHD	TRANSP and Measured Stored Energy (against time)

Table H.1: Inputs to the ‘transpRead.py’ program in order to plot a variety of TRANSP inputs and outputs.

in the user’s directory (at ASDEX Upgrade), one can run the TRANSP post-processing routines to plot the TRANSP inputs and outputs.

First, call ‘python transpRead.py’ or ‘python transpReadAUG2.py.’ The code will ask for a run number. On Alcator C-Mod, this is the TRANSP run number, not the discharge number. On ASDEX Upgrade, this is the discharge number, followed by ‘AXX’, where XX indicates the TRANSP run for that discharge. For example, one would enter ‘33585A04’ for the 04 TRANSP run for discharge 33585. If no run is entered, the code will use the most recent TRANSP run for that discharge.

The code will then ask what the user wants to plot. Enter either a number or

string to select which output to plot. Table H.1 gives the number and string inputs to plot a variety of inputs to and outputs from TRANSP. For quantities that are plotted against radius, the code will also ask for a time or time range over which to average the data.

The most relevant quantities for the validation studies in this thesis are the electron and ion heat fluxes, which are item 10 in Table H.1. Digging into the code itself will also give the reader a sense of how to extract the actual data from the TRANSP run for direct comparison to simulation data.

H.4 GENE

This section does not attempt to explain how to run GENE, as this is a complicated process, and is described in much greater detail in Reference [26]. This section instead focuses on a few post-processing routines that were written as part of this thesis in order to analyze GENE outputs and produce figures for the thesis.

All of these codes are found in the `‘/global/homes/a/acreeley/genecode/python-Routines/’` directory on the NERSC server.

1. Linear Growth Rate Plotter

The Python script `‘plotLinear.py’` requires a `.log` file, which is the output of GENE run linearly in the initial value configuration. It is essentially a text file containing run numbers, wavenumbers, linear growth rates, and real frequencies. Run the code with the following command line call: `python plotLinear.py run_name`. The code will append `‘scan’` to the beginning of `‘run_name’` and `‘.log’` to the end before looking for the file. Entering more than one `‘run_name’` will overplot the growth rate spectra for up to three linear scans.

The code plots the linear growth rate against the wavenumber on a log-log plot and a the real frequency against the wavenumber on a linear-log plot.

2. Dominant and Sub-Dominant Linear Growth Rates

In addition to just plotting the dominant linear growth rates, one can also plot both dominant and subdominant modes using ‘plotEV.py.’ One must run GENE in the eigenvalue configuration in order to calculate both dominant and subdominant mode growth rates. Call: `python plotEV.py IVrun EVrun`, where ‘IVrun’ is the .log file for the initial value run and ‘EVrun’ is the .log file for the eigenvalue run.

The code plots just the dominant linear growth rate against wavenumber, the real frequency of the dominant mode, and then both dominant and subdominant mode linear growth rates against wavenumber on the same plot.

This code was used to generate Figure 5-8.

3. Nonlinear Heat Flux Averaging

GENE generates a huge set of outputs with information about many different aspects of the plasma behavior. In particular, this thesis has been interested in the turbulent heat flux flowing through the ion and electron channels, and a post processing routine was written to calculate this information based on the outputs of GENE, the Python script `nrgReader.py`, run with the following call: `python nrgReader.py nrg_file tStart`.

This code reads an `nrg.dat` file, which should contain the concatenated data from all of the GENE code restarts. The code plots the sum of the electrostatic and electromagnetic heat fluxes for both ions and electrons (in units of the Gyro-Bohm hat flux) against simulation time (in units of a/c_s). One must enter the time at which to start averaging of the data as ‘tStart.’ Generally this time should be chosen such that it is after the linear growth period of the turbulence and some significant time after the effects of $E \times B$ shear have been turned on in the code (if they were not active from the beginning of the simulation). The

code averages the heat flux data from this time until the end of the simulation output.

On the plot, the horizontal dashed lines represent the average heat fluxes. The average value is printed in the terminal, along with an uncertainty, calculated as in Reference [147]. Some additional information about the uncertainty calculation is also printed in the terminal. This procedure was used to calculate all of the turbulent heat fluxes from GENE outputs in this thesis.

This code was used to generate Figure 5-9.

4. Simulation Perturbative Diffusivity Calculation

The Python script ‘geneChiPertPlot33585in.py’ was used to generate Figure 5-12. This script is less interactive than the others described here, but is used to calculate the perturbative diffusivity from the outputs of GENE and compare those to the experimentally measured value. The command line call is simply: `python geneChiPertPlot33585in.py` . All modification must be done inside the script itself. One should change the plasma density, temperature, minor radius, flux surface area, and Gyro-Bohm heat flux unit, as well entering the list of normalized electron temperature gradient scale lengths and corresponding heat fluxes and uncertainties. Running the code then fits the set of points linearly, quadratically, and exponentially, printing the perturbative diffusivity calculated with each method. The code also generates a plot of the results.

H.5 TGLF and VITALS

This section describes the workflow of the novel multi-machine, multi-discharge, multi-constraint validation study with TGLF and VITALS presented in Chapter 6. This workflow will generate, for example, a single point in Figure 6-2 or a single cell in Table 6.3. Details of the data analysis and how to run individual codes are

found either earlier in this Appendix or in other sources. VITALS is an optimization wrapper that iterates many instances of TGLF in order to find the inputs that give the best agreement with the available validation constraints [156]. VITALS is run through OMFIT, a Python wrapper used to perform a variety of transport analysis tasks [157]. Running TGLF and OMFIT requires a number of computer accounts and data access, which are not described in detail here.

1. Select Discharge

First, one must select the discharge (or discharges) and time range (or ranges) that one would like to analyze. While this may seem trivial, one must ensure that each of the measurements required for the validation study are available. For power balance, one requires measurements of: electron density, electron temperature, ion temperature, plasma rotation, effective charge, radiated power, magnetic equilibrium, and heating power (everything that is required to run TRANSP). In practice, the ion temperature and plasma rotation measurements tend to be the most constraining, as they are generally more difficult. On ASDEX Upgrade, they require neutral beam injection to get CXRS measurements, at least in blips for diagnostic purposes. On Alcator C-Mod, they require trace impurity injection to get XICS measurements. More details can be found in Chapter 2

In addition, one must ensure that there are electron temperature fluctuation measurements from CECE in the discharge. This requires that the plasma density be low enough to avoid ECE cutoff at the location of interest. In addition, the fluctuation level must be large enough to be above the sensitivity limit of the CECE diagnostic (this rules out many I- and H-mode plasmas). The magnetic field of the discharge must also be such that the CECE channels are measuring at an appropriate location to get core turbulence measurements. More detail was given in Chapter 4

The discharge must also have partial sawtooth crashes in order to measure the perturbative diffusivity. This requires that the safety factor be below unity on axis, and that the sawtooth period be long enough that partial sawteeth form, instead of just full sawteeth. The discharge should also have a sufficient number of partial sawtooth heat pulses that one can find some reasonable average and uncertainty in the perturbative diffusivity (at least three heat pulses, though preferably more than ten). More detail was given in Chapter 3.

Ideally, one would be able to measure electron temperature fluctuations and the perturbative diffusivity in the same location (as was done for the Alcator C-Mod discharges in Chapter 6), but if not, one can perform validation at two different radial locations (as was done for the ASDEX Upgrade discharges in Chapter 6).

Finally, if one is interested in validation on a series of similar discharges, one must find a series of discharges that all fulfill all of the requirements listed here.

2. Calculate Perturbative Diffusivity

Once one has selected the discharge for analysis, one must calculate the perturbative diffusivity based on the methods presented in Chapter 3 and Section H.1. Running VITALS requires a value of the perturbative diffusivity (in m^2s^{-1}), an uncertainty on this value, and the radial location of the measurement (typically taken as the center of the domain over which the heat pulse is tracked).

3. Process CECE Data

Next, one must analyze the CECE data to extract the electron temperature fluctuation level. For comparisons with TGLF, only the total fluctuation level is used, not the fluctuation spectrum. Details of this analysis are given in Chapter 4 and Section H.2. VITALS requires the total fluctuation level, its uncertainty, and the radial location of the measurement. One must also know

the beam spot size (the $1/e$ electric field diameter) at the measurement location to properly apply a synthetic diagnostic to the TGLF data.

In addition, as discussed in Chapter 6, if one is analyzing data from a low density ASDEX Upgrade discharge, one must correct for the difference between the total temperature fluctuation level (output by TGLF) and the perpendicular temperature fluctuation level (measured by CECE). In practice, this is easiest to correct for by multiplying the experimentally measured value by a correction factor (the fraction $(\tilde{T}_{e,tot}/T_{e,tot})/(\tilde{T}_{e,\perp}/T_{e,\perp})$) before entering the value into VITALS. This correction factor can be estimated from GENE runs on similar plasmas on ASDEX Upgrade.

4. Prepare Data and Run TRANSP

The final preparatory step to running VITALS is to run TRANSP on the discharge and time range of interest in order to get the power balance results. This also produces other inputs that are necessary to run TGLF. Again, details of running TRANSP are found in numerous other publications, but in brief it requires fitting plasma parameter profiles, writing magnetic geometry and other parameter files, generating a TRANSP Namelist file, and then submitting all of these files to the TRANSP servers at PPPL.

5. Import TRANSP Namelist and NetCDF Files into OMFIT

Once all of the experimental data analysis has been completed, one can begin setting up the validation work within OMFIT. First, one should import the TRXPL module into OMFIT. Within the TRXPL module, load the TRANSP Namelist and the NetCDF file from the TRANSP run into OMFIT. For Alcator C-Mod this can be done directly from the MDSplus server. For ASDEX Upgrade, one must manually import the TRANSP Namelist file (typically named, for example, ‘33585TR.DAT’) and NetCDF file (typically named, for example,

'33585.cdf') from the user's file directory on the ASDEX Upgrade computers. One should delete the TRANSP run number if it is included in these file names (i.e. the file name should be '33585TR.DAT' and not '33585A04TR.DAT').

6. Run TRXPL to generate a Statefile and g-File

Once the inputs from TRANSP have been loaded into TRXPL, one should run TRXPL in order to generate a Statefile and a g-File. The Statefile contains the plasma profiles and other information about the discharge. The g-File contain information about the magnetic geometry of the discharge. These are required inputs for VITALS and TGLF.

7. Import Statefile and g-File into VITALS and Begin Pre-Processing

Now, import the VITALS module into OMFIT. Open the VITALS module and import the Statefile and g-File from TRXPL into VITALS. Once this is complete, one can use the VITALS GUI to generate input profiles (using PROFILES_GEN) and inputs to TGLF. It is at this point that one chooses whether TGLF will run with the standard k-spectrum (multi-scale) or with a user-defined k-spectrum extending only up to $k_y \rho_s = 1.0$ (ion-scale). One also chooses a radius at which to run the simulations. At this point, all of the data pre-processing is complete, and one only has to input the settings and constraints into VITALS.

8. Input Validation Constraints and Select Code Inputs to Vary

One must now decide which validation constraints one will use to constrain the TGLF simulations. By default, one should probably include the ion and electron heat fluxes. For these constraints, one must only input an uncertainty, as VITALS will automatically read the outputs of TRANSP in order to retrieve the experimental value of the heat fluxes. TRANSP uncertainty calculations are described in Chapter 2. To use the perturbative diffusivity as a validation con-

straint, one must input both the value of the measured perturbative diffusivity and the uncertainty on this value.

Finally, to include the CECE electron temperature fluctuation measurements as a constraint, one must include the total fluctuation level, an uncertainty on this level, and the value of $k_y \rho_s$ that corresponds to the beam spot size. To calculate this, evaluate $(2\pi/w_0)\rho_s$ where w_0 is $1/e$ electric field diameter and ρ_s is the ion Larmor radius evaluated at the electron temperature. If ρ_s is not known already, TGLF will calculate this during its first run. One can therefore run VITALS once without CECE as a constraint, extract the value of ρ_s , calculate the expression above, input this into the CECE section, and then rerun VITALS with CECE as a constraint.

Finally, one must choose which code inputs to vary within uncertainty, and provide the experimental uncertainty for each of these parameters. Typically this includes at least the electron and ion temperature gradients. For the study in Chapter 6 this also included the effective charge Z_{eff} and the electron density gradient. Uncertainties were calculated using experimental and profile fitting uncertainties.

9. Run VITALS

After all of the inputs are set, one is ready to run VITALS. The code will iterate many runs of TGLF, varying the inputs within their uncertainty in order to optimally match the set of validation constraints. The code will eventually output a final run, giving the best set of input parameters and the corresponding closest match of all of the validation constraints. Note that in some cases, VITALS may optimize to have one constraint above one standard deviation from the measurement and others lower, as opposed to having all of them just below one standard deviation. One can look at earlier iterations of the code

to search for these cases. If one is interested in the linear growth spectrum of a particular run of TGLF, one can also find this output in the temporary storage of the VITALS run. One can also download this spectrum to OMFIT. In general, this thesis used the spectrum for the nominal experimental inputs to calculate the criteria for the importance of multi-scale turbulent interactions.

10. Calculate Ricci Metric and Criteria

Using the output of the closest matches of all of the validation constraints for a particular experimental case, one can then calculate the Ricci Validation Metric, as described in Appendix D. One can then calculate the values of the two criteria for the importance of multi-scale effects, using the linear growth spectra from the output of TGLF.

Finally, if one returns to Step 7 and changes the code from multi-scale to ion-scale, and then reruns VITALS, one can calculate the difference in the Ricci metric between the ion- and multi-scale runs, thus calculating a single point on Figure 6-2 or Figure 6-4. Other cells in a given row in Table 6.3 involve changing which inputs are varied within uncertainty and which validation constraints are applied to a particular plasma condition. One then repeats this entire process for a different discharge in order to calculate another point in the figures or row in the table.

Bibliography

- [1] AIAA, “Guide for the verification and validation of computational fluid dynamics simulations,” Tech. Rep. AIAA G-077-1998(2002), The American Institute of Aeronautics and Astronautics, 1998.
- [2] C. M. Surko and R. E. Slusher *Science*, vol. 221, p. 817, 1983.
- [3] J. W. Conner and H. R. Wilson *Plasma Phys. Control. Fusion*, vol. 36, p. 719, 1994.
- [4] W. Horton *Rev. Mod. Phys.*, vol. 71 (3), p. 735, 1999.
- [5] G. R. Tynan, A. Fujisawa, and G. McKee *Plasma Phys. Control. Fusion*, vol. 51, p. 113001, 2009.
- [6] J. Freidberg, *Plasma Physics and Fusion Energy*. Cambridge University Press, 2007.
- [7] C. Sung. Ph.D. thesis, Massachusetts Institute of Technology, 2014.
- [8] F. F. Chen, *Introduction to Plasma Physics and Controlled Fusion*. Springer, 1984.
- [9] F. Jenko, W. Dorland, and M. Kotschenreuther *Phys. Plasmas*, vol. 7, p. 1904, 2000.
- [10] N. T. Howard, C. Holland, A. E. White, M. Greenwald, and J. Candy *Nucl. Fusion*, vol. 56, p. 014004, 2016.
- [11] N. T. Howard, A. E. White, M. L. Reinke, M. Greenwald, C. Holland, J. Candy, and J. R. Walk *Nucl. Fusion*, vol. 53, p. 123011, 2013.
- [12] N. T. Howard, C. Holland, A. E. White, M. Greenwald, and J. Candy *Phys. Plasmas*, vol. 21, p. 112510, 2014.
- [13] S. Maeyama, Y. Idomura, T.-H. Watanabe, M. Nakata, M. Yagi, N. Miyato, A. Ishizawa, and M. Nunami *Phys. Rev. Lett.*, vol. 114, p. 255002, 2015.
- [14] N. T. Howard, C. Holland, A. E. White, M. Greenwald, J. Candy, and A. J. Creely *Phys. Plasmas*, vol. 23, p. 056109, 2016.

- [15] C. Holland, N. T. Howard, and B. A. Grierson *Nucl. Fusion*, vol. 57, p. 066043, 2017.
- [16] S. Maeyama, T.-H. Watanabe, and A. Ishizawa *Phys. Rev. Lett.*, vol. 119, p. 195002, 2017.
- [17] N. T. Howard, C. Holland, A. E. White, M. Greenwald, P. Rodriguez-Fernandez, J. Candy, and A. J. Creely *Plasma Phys. Control. Fusion*, vol. 60, p. 014034, 2018.
- [18] J. Candy and R. E. Waltz *J. Comput. Phys.*, vol. 186, pp. 545–581, 2003.
- [19] G. M. Staebler, J. E. Kinsey, and R. E. Waltz *Phys. Plasmas*, vol. 12, p. 102508, 2005.
- [20] G. M. Staebler, J. E. Kinsey, and R. E. Waltz *Phys. Plasmas*, vol. 14, p. 055909, 2007.
- [21] G. M. Staebler, J. Candy, N. T. Howard, and C. Holland *Phys. Plasmas*, vol. 23, p. 062518, 2016.
- [22] G. M. Staebler, N. T. Howard, J. Candy, and C. Holland *Nucl. Fusion*, vol. 57, p. 066046, 2017.
- [23] X. Garbet, Y. Idomura, L. Villard, and T. H. Watanabe *Nucl. Fusion*, vol. 50, p. 043002, 2010.
- [24] W. Dorland and G. W. Hammett *Phys. of Fluids B*, vol. 5, p. 812, 1993.
- [25] J. Candy, E. A. Belli, and R. V. Bravenec *J. Comput. Phys.*, vol. 324, p. 73, 2016.
- [26] Gene Development Team, *The Gyrokinetic Plasma Turbulence Code Gene: User Manual*, 2015.
- [27] R. L. Miller, M. S. Chu, J. M. Greene, Y. R. Lin-Liu, and R. E. Waltz *Phys. Plasmas*, vol. 5, p. 973, 1998.
- [28] P. Mantica, C. Angioni, C. Challis, G. Colyer, L. Frassinetti, N. Hawkes, T. Johnson, M. Tsalas, P. C. deVries, J. Weiland, B. Baiocchi, M. N. A. Beurskens, A. C. A. Figueiredo, C. Giroud, J. Hobirk, E. Joffrin, E. Lerche, V. Naulin, A. G. Peeters, A. Salmi, C. Sozzi, D. Srintzi, G. Staebler, T. Tala, D. Van Eester, and T. Versloot *Phys. Rev. Lett.*, vol. 107, p. 135004, 2011.
- [29] J. Citrin, F. Jenko, P. Mantica, D. Told, C. Bourdelle, J. Garcia, J. W. Haverkort, G. M. D. Hogeweij, T. Johnson, and M. J. Pueschel *Phys. Rev. Lett.*, vol. 111, p. 155001, 2013.
- [30] C. Bourdelle, X. Garbet, F. Imbeaux, A. Casati, N. Dubuit, R. Guirlet, and T. Parisot *Phys. Plasmas*, vol. 14, p. 112501, 2007.

- [31] M. Greenwald, A. Bader, S. Baek, M. Bakhtiari, H. Barnard, W. Beck, W. Bergerson, I. Bespamyatnov, P. Bonoli, D. Brower, D. Brunner, W. Burke, J. Candy, M. Churchill, I. Cziegler, A. Diallo, A. Dominguez, B. Duval, E. Edlund, P. Ennever, D. Ernst, I. Faust, C. Fiore, T. Fredian, O. Garcia, C. Gao, J. Goetz, T. Golfinopoulos, R. Granetz, O. Grulke, Z. Hartwig, S. Horne, N. Howard, A. E. Hubbard, J. Hughes, I. Hutchinson, J. Irby, V. Izzo, C. Kessel, B. LaBombard, C. Lau, C. Li, Y. Lin, B. Lipschultz, A. Loarte, E. Marmar, A. Mazurenko, G. McCracken, R. McDermott, O. Meneghini, D. Mikkelsen, D. Mossessian, R. Mumgaard, J. Myra, E. Nelson-Melby, R. Ochoukov, G. Olynyk, R. Parker, S. Pitcher, Y. Podpaly, M. Porkolab, M. Reinke, J. Rice, W. Rowan, A. Schmidt, S. Scott, S. Shiraiwa, J. Sierchio, N. Smick, J. A. Snipes, P. Snyder, B. Sorbom, J. Stillerman, C. Sung, Y. Takase, V. Tang, J. Terry, D. Terry, C. Theiler, A. Tronchin-James, N. Tsujii, R. Vieira, J. Walk, G. Wallace, A. E. White, D. G. Whyte, J. Wilson, S. Wolfe, G. Wright, J. Wright, S. Wukitch, and S. Zweben *Phys. Plasmas*, vol. 21, p. 110501, 2014.
- [32] J. W. Hughes, D. Mossessian, K. Zhurovich, M. DeMaria, K. Jensen, and A. Hubbard *Rev. Sci. Instrum.*, vol. 74, p. 1667, 2003.
- [33] J. O'Shea, A. E. Hubbard, and the Alcator C-Mod Group in *9th Joint Workshop on ECE and ECRH Borrego Springs*, p. 7, 1995.
- [34] P. J. L. O'Shea. Ph.D. thesis, Massachusetts Institute of Technology, 1997.
- [35] J. H. Irby, E. S. Marmar, E. Sevillano, and S. M. Wolfe *Rev. Sci. Instrum.*, vol. 59, p. 1568, 1988.
- [36] A. Ince-Cushman, J. E. Rice, M. Bitter, M. L. Reinke, K. W. Hill, M. F. Gu, E. Eikenberry, C. Broennimann, S. Scott, Y. Podpaly, S. G. Lee, and E. S. Marmar *Rev. Sci. Instrum.*, vol. 79, p. 10E302, 2008.
- [37] M. E. Foord, E. S. Marmar, and J. L. Terry *Rev. Sci. Instrum.*, vol. 53 (9), p. 1407, 1982.
- [38] J. A. Goetz, B. Lipschultz, M. A. Graf, C. Kurz, R. Nachtrieb, J. A. Snipes, and J. L. Terry *J. Nucl. Mater.*, vol. 220, p. 971, 1995.
- [39] L. L. Lao, H. St. John, R. D. Stambaugh, A. G. Kellman, and W. Pfeiffer *Nucl. Fusion*, vol. 25, p. 1611, 1985.
- [40] A. Kallenbach, the ASDEX Upgrade Team, and the EUROfusion MST1 Team *Nucl. Fusion*, vol. 57, p. 102015, 2017.
- [41] H. Murmann, S. Goetsch, H. Roehr, H. Salzmann, and K. H. Steuer *Rev. Sci. Instrum.*, vol. 63, p. 4941, 1992.
- [42] B. Kurzan and H. D. Murmann *Rev. Sci. Instrum.*, vol. 82, p. 103501, 2011.

- [43] M. Willensdorfer, S. S. Denk, E. Strumberger, W. Suttrop, B. Vanovac, D. Brida, M. Cavedon, I. Classen, M. Dunne, S. Fietz, R. Fischer, A. Kirk, F. M. Lagner, Y. Q. Liu, T. Odstrcil, D. A. Ryan, E. Viezzer, H. Zohm, I. C. Luhmann, the ASDEX Upgrade Team, and the EUROfusion MST1 Team *Plasma Phys. Control. Fusion*, vol. 58, p. 114004, 2016.
- [44] A. Mlynek, G. Schramm, H. Eixenberger, G. Sips, K. McCormick, M. Zilker, K. Behler, J. Eheberg, and the ASDEX Upgrade Team *Rev. Sci. Instrum.*, vol. 81, p. 033507, 2010.
- [45] E. Viezzer, T. Puetterich, R. Dux, R. M. McDermott, and the ASDEX Upgrade Team *Rev. Sci. Instrum.*, vol. 83, p. 103501, 2012.
- [46] H. Meister, R. Fischer, L. D. Horton, C. F. Maggi, D. Nishijima, the ASDEX Upgrade Team, C. Giroud, K. D. Zastrow, JET-EFDA Contributors, and B. Zaniol *Rev. Sci. Instrum.*, vol. 75, p. 4097, 2004.
- [47] K. F. Mast, J. C. Vallet, C. Andelfinger, P. Betzler, H. Kraus, and G. Schramm *Rev. Sci. Instrum.*, vol. 62, p. 744, 1991.
- [48] P. J. McCarthy *Phys. Plasmas*, vol. 6, p. 3554, 1999.
- [49] C. Holland *Phys. Plasmas*, vol. 23, p. 060901, 2016.
- [50] M. A. Chilenski. Ph.D. thesis, Massachusetts Institute of Technology, 2016.
- [51] J. M. Schachter. Ph.D. thesis, Massachusetts Institute of Technology, 1997.
- [52] See <http://w3.pppl.gov/transp> for full documentation concerning the TRANSP code.
- [53] N. J. L. Cardozo *Plasma Phys. Control. Fusion*, vol. 37, p. 799, 1995.
- [54] B. J. D. Tubbing, N. J. L. Cardozo, and M. J. Van der Wiel *Nucl. Fusion*, vol. 27, p. 1843, 1987.
- [55] A. J. Creely, A. E. White, E. M. Edlund, N. T. Howard, and A. E. Hubbard *Nucl. Fusion*, vol. 56, p. 036003, 2016.
- [56] F. Ryter, R. Dux, P. Mantica, and T. Tala *Plasma Phys. Control. Fusion*, vol. 52, p. 124043, 2010.
- [57] J. D. Callen and G. L. Jahns *Phys. Rev. Lett.*, vol. 38, p. 491, 1977.
- [58] E. D. Fredrickson, M. E. Austin, R. Groebner, J. Manickam, B. Rice, G. Schmidt, and R. Snider *Phys. Plasmas*, vol. 7, p. 5051, 2000.
- [59] F. Ryter, F. Leuterer, G. Pereverzev, H.-U. Fahrback, J. Stober, W. Suttrop, and the ASDEX Upgrade Team *Phys. Rev. Lett.*, vol. 86, p. 2325, 2001.

- [60] P. Mantica, F. D. Luca, G. Gorini, and A. Jacchia *Nucl. Fusion*, vol. 32, p. 2203, 1992.
- [61] N. J. L. Cardozo and J. C. M. de Haas *Nucl. Fusion*, vol. 30, p. 521, 1990.
- [62] N. J. L. Cardozo and A. C. C. Sips *Plasma Phys. Control. Fusion*, vol. 33, p. 1337, 1991.
- [63] S. P. Smith, C. C. Petty, A. E. White, C. Holland, R. Bravenec, M. E. Austin, L. Zeng, and O. Meneghini *Nucl. Fusion*, vol. 55, p. 083011, 2015.
- [64] A. J. Creely, T. Görler, G. D. Conway, S. J. Freethy, N. T. Howard, P. A. Schneider, A. E. White, M. Willensdorfer, and the ASDEX Upgrade Team *Nucl. Fusion*, vol. 58, p. 126001, 2018.
- [65] A. J. Creely, P. Rodriguez-Fernandez, G. D. Conway, S. J. Freethy, N. T. Howard, A. E. White, and the ASDEX Upgrade Team *Plasma Phys. Control. Fusion*, vol. Submitted, p. NA, 2019.
- [66] I. H. Hutchinson, *Principles of Plasma Diagnostics*. IOP Publishing, 2002.
- [67] H. Hartfuss and T. Geist, *Fusion Plasma Diagnostics with mm-Waves*. Wiley, 2013.
- [68] M. Bornatici, R. Cano, O. de Barbieri, and F. Engelmann *Nucl. Fusion*, vol. 23, p. 1153, 1983.
- [69] S. Sattler and H. J. Hartfuss *Phys. Rev. Lett.*, vol. 72, p. 653, 1993.
- [70] C. Watts, R. F. Gandy, T. D. Rempel, and G. Cima *Rev. Sci. Instrum.*, vol. 66, p. 451, 1995.
- [71] G. Cima, R. V. Bravenec, A. J. Wootton, T. D. Rempel, R. F. Gandy, C. Watts, and M. Kwon *Phys. Plasmas*, vol. 2, p. 720, 1995.
- [72] C. Watts, Y. In, J. Heard, P. Phillips, A. Lynn, A. Hubbard, and R. Gandy *Nucl. Fusion*, vol. 44 (9), p. 987, 2004.
- [73] C. Watts *Fusion Sci. Technol.*, vol. 52, p. 176, 2007.
- [74] A. E. White, L. Schmitz, W. A. Peebles, T. A. Carter, T. L. Rhodes, E. J. Doyle, P. A. Gourdain, J. C. Hillesheim, G. Wang, C. Holland, G. R. Tynan, M. E. Austin, G. R. McKee, M. W. Shafer, K. H. Burrell, J. Candy, J. C. DeBoo, R. Prater, G. M. Staebler, R. E. Waltz, and M. A. Makowski *Rev. Sci. Instrum.*, vol. 79 (10), p. 103505, 2008.
- [75] C. Sung, A. E. White, J. H. Irby, R. Leccacorvi, R. Vieira, C. Y. Oi, W. A. Peebles, and X. Nguyen *Rev. Sci. Instrum.*, vol. 83, p. 10E311, 2012.

- [76] N. T. Howard, C. Sung, and A. E. White *Rev. Sci. Instrum.*, vol. 85, p. 11D811, 2014.
- [77] S. E. Zemedkun, S. Che, Y. Chen, C. W. Domier, N. C. Luhmann Jr., T. Mun-sat, S. E. Parker, B. Tobias, W. Wan, and L. Yu *Phys. Plasmas*, vol. 22, p. 122508, 2015.
- [78] C. Sung, W. A. Peebles, C. Wannberg, T. L. Rhodes, X. Nguyen, R. Lantsov, and L. BardÚczi *Rev. Sci. Instrum.*, vol. 87, p. 11E123, 2016.
- [79] M. Fontana, L. Porte, and P. M. Cabrera *Rev. Sci. Instrum.*, vol. 88, p. 083506, 2017.
- [80] S. J. Freethy, G. D. Conway, I. Classen, A. J. Creely, T. Happel, A. KŽhn, B. Vanovac, and A. E. White *Rev. Sci. Instrum.*, vol. 87, p. 11E102, 2016.
- [81] A. J. Creely, S. J. Freethy, W. M. Burke, G. D. Conway, R. Leccacorvi, W. C. Parkin, D. R. Terry, and A. E. White *Rev. Sci. Instrum.*, vol. 89, p. 053503, 2018.
- [82] J. E. Rice, C. Gao, M. L. Reinke, P. H. Diamond, N. T. Howard, H. J. Sun, I. Cziegler, A. E. Hubbard, Y. A. Podpaly, W. L. Rowan, J. L. Terry, M. A. Chilenski, L. Delgado-Aparicio, P. C. Ennever, D. Ernst, M. J. Greenwald, J. W. Hughes, Y. Ma, E. S. Marmor, M. Porkolab, A. E. White, and S. M. Wolfe *Nucl. Fusion*, vol. 53, p. 033004, 2013.
- [83] V. S. Udintsev, M. Ottaviani, P. Maget, G. Giruzzi, J.-L. Segui, T. Aniel, J. F. Artaud, F. Clairet, M. Goniche, G. T. Hoang, G. T. A. Huysmans, F. Imbeaux, E. Joffrin, D. Mazon, A. L. Pecquet, R. Sabot, A. Sirinelli, L. Vermare, the Tore Supra Team, A. Kramer-Flecken, H. R. Koslowski, the TEXTOR Team, A. J. H. Donne, F. C. Schuller, C. W. Domier, N. C. Luhmann Jr., and S. V. Mirnov *Plasma Phys. Control. Fusion*, vol. 47, p. 1111, 2005.
- [84] N. J. L. Cardozo, B. J. D. Tubbing, F. Tibone, and A. Taroni *Nucl. Fusion*, vol. 28, p. 1173, 1988.
- [85] A. Jacchia, P. Mantica, F. D. Luca, and G. Gorini *Phys. Fluids B*, vol. 3, p. 11, 1991.
- [86] F. Ryter, G. Tardini, F. D. Luca, H. U. Fahrbach, F. Imbeaux, A. Jacchia, K. K. Kirov, F. Leuterer, P. Mantica, A. G. Peeters, G. Pereverzev, W. Suttrop, and the ASDEX Upgrade Team *Nucl. Fusion*, vol. 43, p. 1396, 2003.
- [87] F. Ryter, C. Angioni, C. Giroud, A. G. Peeters, T. Biewer, R. Bilato, E. Joffrin, T. Johnson, H. Leggate, E. Lerche, G. Madison, P. Mantica, D. Van Eester, I. Voitsekhovitch, and JET Contributors *Nucl. Fusion*, vol. 51, p. 113016, 2011.

- [88] P. A. Schneider, A. Bustos, P. Hennequin, F. Ryter, M. Bernert, M. Cave-
don, M. G. Dunne, R. Fischer, T. Görler, T. Happel, V. Igochine, B. Kurzan,
A. Lebschy, R. M. McDermott, P. Morel, M. Willensdorfer, the ASDEX Up-
grade Team, and the EUROfusion MST1 Team *Nucl. Fusion*, vol. 57, p. 066003,
2017.
- [89] D. G. Whyte, A. E. Hubbard, J. W. Hughes, B. Lipschultz, J. E. Rice, E. S. Mar-
mar, M. Greenwald, I. Cziegler, A. Dominguez, T. Golfopoulos, N. Howard,
L. Lin, R. M. McDermott, M. Porkolab, M. L. Reinke, J. Terry, N. Tsujii,
S. Wolfe, S. Wukitch, Y. Lin, and the Alcator C-Mod Team *Nucl. Fusion*,
vol. 50, p. 105005, 2010.
- [90] A. E. Hubbard, D. G. Whyte, R. M. Churchill, A. Dominguez, J. W. Hughes,
Y. Ma, E. S. Marmor, Y. Lin, M. L. Reinke, and A. E. White *Nucl. Fusion*,
vol. 52, 2012.
- [91] A. E. Hubbard, T. Osborne, F. Ryter, M. Austin, L. B. Orte, R. M. Churchill,
I. Cziegler, M. Fenstermacher, R. Fischer, S. Gerhardt, R. Groebner, P. Gohil,
T. Happel, J. W. Hughes, A. Loarte, R. Maingi, P. Manz, A. Marinoni, E. S.
Marmor, R. M. McDermott, G. McKee, T. L. Rhodes, J. E. Rice, L. Schmitz,
C. Theiler, E. Viezzer, J. R. Walk, A. E. White, D. Whyte, S. Wolfe, E. Wol-
frum, Z. Yan, the Alcator C-Mod Team, the ASDEX Upgrade Team, and the
DIII-D Team *Nucl. Fusion*, vol. 56, 2016.
- [92] F. Ryter, R. Fischer, J. C. Fuchs, T. Happel, R. M. McDermott, E. Viezzer,
E. Wolfrum, L. B. Orte, M. Bernert, A. Burckhart, S. da Graca, B. Kurzan,
P. McCarthy, T. Pütterich, W. Suttrop, M. Willensdorfer, and the ASDEX
Upgrade Team *Nucl. Fusion*, vol. 57, p. 016004, 2017.
- [93] A. Marinoni, J. C. Rost, M. Porkolab, A. E. Hubbard, T. H. Osborne, A. E.
White, D. G. Whyte, T. L. Rhodes, E. M. Davis, D. R. Ernst, and K. H. Burrell
Nucl. Fusion, vol. 55, p. 093019, 2015.
- [94] A. E. Hubbard, S. G. Baek, D. Brunner, A. J. Creely, I. Cziegler, E. Edlund,
J. W. Hughes, B. LaBombard, Y. Lin, Z. Liu, E. S. Marmor, M. L. Reinke,
J. E. Rice, B. Sorbom, C. Sung, J. Terry, C. Theiler, E. A. Tolman, J. R. Walk,
A. E. White, D. Whyte, S. M. Wolfe, S. Wukitch, X. Q. Xu, and the Alcator
C-Mod Team *Nucl. Fusion*, vol. 57, 2017.
- [95] See <https://fusion.gat.com/theory/Gyrodock> for full documentation concerning
the GYRO code.
- [96] A. E. White, M. Barnes, A. Dominguez, M. Greenwald, N. T. Howard, A. E.
Hubbard, J. W. Hughes, D. R. Mikkelsen, F. I. Parra, M. L. Reinke, C. Sung,
J. Walk, and D. G. Whyte *Nucl. Fusion*, vol. 54, p. 083019, 2014.
- [97] A. E. White, N. T. Howard, A. J. Creely, M. A. Chilenski, M. Greenwald, A. E.
Hubbard, J. W. Hughes, E. Marmor, J. E. Rice, J. M. Sierchio, C. Sung, J. R.

- Walk, D. G. Whyte, D. R. Mikkelsen, E. M. Edlund, C. Kung, C. Holland, J. Candy, C. C. Petty, M. L. Reinke, and C. Theiler *Phys. Plasmas*, vol. 22, p. 056109, 2015.
- [98] See <http://www.mathworks.com/products/matlab/> for documentation concerning MATLAB and its functions.
- [99] X. Garbet, P. Mantica, F. Ryter, G. Cordey, F. Imbeaux, C. Sozzi, A. Manini, E. Asp, V. Parail, R. Wolf, and JET EFDA Contributors *Plasma Phys. Control. Fusion*, vol. 46, p. 1351, 2004.
- [100] A. J. Creely, N. T. Howard, P. Rodriguez-Fernandez, N. Cao, A. E. Hubbard, J. W. Hughes, J. E. Rice, A. E. White, J. Candy, G. M. Staebler, G. D. Conway, S. J. Freethy, and C. Sung *Phys. Plasmas*, vol. 24, p. 056104, 2017.
- [101] F. Ryter, C. Angioni, M. Beurskens, S. Cirant, G. T. Hoang, G. M. D. Hogeweij, F. Imbeaux, A. Jacchia, P. Mantica, W. Suttrop, and G. Tardini *Plasma Phys. Control. Fusion*, vol. 43, p. A323, 2001.
- [102] D. Told, F. Jenko, T. Görler, F. J. Casson, E. Fable, and the ASDEX Upgrade Team *Phys. Plasmas*, vol. 20, p. 122312, 2013.
- [103] C. Angioni, A. G. Peeters, F. Jenko, and T. Dannert *Phys. Plasmas*, vol. 12, p. 112310, 2005.
- [104] F. Ryter, C. Angioni, A. G. Peeters, F. Leuterer, H. U. Fahrbach, W. Suttrop, and the ASDEX Upgrade Team *Phys. Rev. Lett.*, vol. 95, p. 085001, 2005.
- [105] C. Sung, A. E. White, N. T. Howard, C. Y. Oi, J. E. Rice, C. Gao, P. Ennever, M. Porkolab, F. Parra, D. Mikkelsen, D. Ernst, J. Walk, J. W. Hughes, J. Irby, C. Kasten, A. E. Hubbard, M. J. Greenwald, and the Alcator C-Mod Team *Nucl. Fusion*, vol. 53, p. 083010, 2013.
- [106] T. Happel, G. D. Conway, W. Kasperek, B. Plaum, C. Lechte, D. Wagner, U. Stroth, and the ASDEX Upgrade Team in *Proceedings of the 10th International Reflectometry Workshop*, pp. 1–6, 2011.
- [107] S. J. Freethy, T. Görler, A. J. Creely, G. D. Conway, T. Happel, C. Koenen, P. Hennequin, A. E. White, and the ASDEX Upgrade Team *Phys. Plasmas*, vol. 25, p. 055903, 2018.
- [108] N. Weiner *Acta Mathematica*, vol. 55, no. 2, p. 117, 1930.
- [109] D. E. Smith, E. J. Powers, and G. S. Caldwell *IEEE Transactions on Plasma Science*, vol. PS-2, pp. 261–272, 1974.
- [110] G. C. Carter and C. H. Knapp *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-23, no. 3, p. 257, 1975.

- [111] K. D. Saunders and F. C. Hanrick *Journal of Geophysical Research*, vol. 87, no. C12, p. 9699, 1982.
- [112] S. L. Marple, *Digital Spectral Analysis: With Applications*. Prentice Hall, 1987.
- [113] G. C. Carter *Proceedings of the IEEE*, vol. 75, no. 2, p. 236, 1987.
- [114] W. H. . Carter, “Coherence theory,” in *Handbook of Optics: Volume I: Fundamentals, Techniques, and Design* (M. Bass, E. W. Van Stryland, D. R. Williams, and W. L. Wolfe, eds.), ch. 4, p. 4.12, McGraw-Hill, Inc., 1995.
- [115] J. S. Bendat and A. G. Piersol, *Random Data*. Wiley, 2 ed., 2000.
- [116] P. Stoica and R. Moses, *Spectral Analysis of Signals*. Prentice Hall, 2002.
- [117] R. Shiavi, *Introduction to Applied Statistical Signal Analysis*. Elsevier Inc., 2007.
- [118] G. Bekefi, *Radiation Processes in Plasmas*. Wiley, 1966.
- [119] C. Watts, H. J. Hartfuss, and M. Häse *Rev. Sci. Instrum.*, vol. 75 (10), p. 3177, 2004.
- [120] S. Sattler and H. J. Hartfuss *Plasma Phys. Control. Fusion*, vol. 35, p. 1285, 1993.
- [121] V. S. Udintsev, M. Goniche, J. L. Segui, G. Giruzzi, D. Molina, F. Turco, G. T. A. Huysmans, P. Maget, the Tore Supra Team, and A. Kraemer-Flecken *Fusion Sci. Technol.*, vol. 50, p. 508, 2006.
- [122] J. R. Walk, J. W. Hughes, A. E. Hubbard, J. L. Terry, D. G. Whyte, A. E. White, S. G. Baek, M. L. Reinke, C. Theiler, R. M. Churchill, J. E. Rice, P. B. Snyder, T. Osborne, A. Dominguez, and I. Cziegler *Phys. Plasmas*, vol. 21, p. 056103, 2014.
- [123] A. E. Hubbard, D. G. Whyte, R. M. Churchill, I. Cziegler, A. Dominguez, T. Golfinopoulos, J. W. Hughes, J. E. Rice, I. Bespamyatnov, M. J. Greenwald, N. Howard, B. Lipschultz, E. S. Marmor, M. L. Reinke, W. L. Rowan, J. L. Terry, and the Alcator C-Mod Team *Phys. Plasmas*, vol. 18, p. 056115, 2011.
- [124] J. E. Rice, M. L. Reinke, C. Gao, N. T. Howard, M. A. Chilenski, L. Delgado-Aparicio, R. S. Granetz, M. J. Greenwald, A. E. Hubbard, J. W. Hughes, J. H. Irby, Y. Lin, E. S. Marmor, R. T. Mumgaard, S. D. Scott, J. L. Terry, J. R. Walk, A. E. White, D. G. Whyte, S. M. Wolfe, and S. J. Wukitch *Nucl. Fusion*, vol. 55, p. 033014, 2015.
- [125] T. Happel, P. Manz, F. Ryter, P. Hennequin, A. Hetzenecker, G. D. Conway, L. Guimaraes, C. Honoré, U. Stroth, E. Viezzer, and the ASDEX Upgrade Team *Nucl. Fusion*, vol. 56, p. 064004, 2016.

- [126] N. T. Howard, A. E. White, M. Greenwald, C. Holland, and J. Candy *Phys. Plasmas*, vol. 21, p. 032308, 2014.
- [127] L. Lin, M. Pokerlab, E. M. Edmund, M. Greenwald, N. Tsujii, J. Candy, R. E. Waltz, and D. R. Mikkelsen *Plasma Phys. Control. Fusion*, vol. 51, p. 065006, 2009.
- [128] L. Lin, M. Porkolab, E. M. Edlund, J. C. Rost, C. L. Fiore, M. Greenwald, Y. Lin, D. R. Mikkelsen, N. Tsujii, and S. J. Wukitch *Phys. Plasmas*, vol. 16, p. 012502, 2009.
- [129] N. T. Howard, M. Greenwald, D. R. Mikkelsen, M. L. Reinke, A. E. White, D. Ernst, Y. Podpaly, and J. Candy *Nucl. Fusion*, vol. 52, p. 063002, 2012.
- [130] R. V. Bravenec and W. M. Nevins *Rev. Sci. Instrum.*, vol. 77, p. 015101, 2006.
- [131] C. Holland, A. E. White, G. R. McKee, M. W. Shafer, J. Candy, R. E. Waltz, L. Schmitz, and G. R. Tynan *Phys. Plasmas*, vol. 16, p. 052301, 2008.
- [132] D. R. Mikkelsen, N. T. Howard, A. E. White, and A. J. Creely *Phys. Plasmas*, vol. 25, p. 042505, 2018.
- [133] A. M. Dimits, G. Bateman, M. A. Beer, B. I. Cohen, W. Dorland, G. W. Hammett, C. Kim, J. E. Kinsey, M. Kotschenreuther, A. H. Kritiz, L. L. Lao, J. Mandrekas, W. M. Nevins, S. E. Parker, A. J. Redd, D. E. Shumaker, R. Sydora, and J. Weiland *Phys. Plasmas*, vol. 7, p. 969, 2000.
- [134] J. Candy, R. Waltz, and W. Dorland *Phys. Plasmas*, vol. 11, p. L25, 2004.
- [135] W. M. Nevins, J. Candy, S. Cowley, T. Dannert, A. Dimits, W. Dorland, C. Estrada-Mila, G. W. Hammett, F. Jenko, M. J. Pueschel, and D. E. Shumaker *Phys. Plasmas*, vol. 13, p. 122306, 2006.
- [136] W. M. Nevins, S. E. Parker, Y. Chen, J. Candy, A. Dimits, W. Dorland, G. W. Hammett, and F. Jenko *Phys. Plasmas*, vol. 14, p. 084501, 2007.
- [137] C. M. Roach, I. G. Abel, R. J. Akers, W. Arter, M. Barnes, Y. Camenen, F. J. Casson, G. Colyer, J. W. Connor, S. C. Cowley, D. Dickinson, W. Dorland, A. R. Field, W. Guttenfelder, G. W. Hammett, R. J. Hastie, E. Highcock, N. F. Loureiro, A. G. Peeters, M. Reshko, S. Saarelma, A. A. Schekochihin, M. Valovic, and H. R. Wilson *Plasma Phys. Control. Fusion*, vol. 51, p. 124020, 2009.
- [138] R. V. Bravenec, C. Holland, J. Candy, and M. Barnes *Phys. Plasmas*, vol. 18, p. 122505, 2011.
- [139] T. Görler, A. E. White, D. Told, F. Jenko, C. Holland, and T. L. Rhodes *Phys. Plasmas*, vol. 21, p. 122307, 2014.

- [140] R. Buchholz, S. Grosshauser, W. Guttenfelder, W. A. Hornsby, P. Migliano, A. G. Peeters, and D. Strintzi *Phys. Plasmas*, vol. 22, p. 082307, 2015.
- [141] R. Bravenec, J. Citrin, J. Candy, P. Mantica, T. Görler, and JET Contributors *Plasma Phys. Control. Fusion*, vol. 58, p. 125018, 2016.
- [142] T. Görler, N. Tronko, W. A. Hornsby, A. Bottino, R. Kleiber, C. Norcini, V. Grandgirard, F. Jenko, and E. Sonnendruecker *Phys. Plasmas*, vol. 23, p. 072503, 2016.
- [143] N. Tronko, A. Bottino, T. Görler, E. Sonnendruecker, D. Told, and L. Villard *Phys. Plasmas*, vol. 24, p. 056115, 2017.
- [144] T. Görler and F. Jenko *Phys. Rev. Lett.*, vol. 100, p. 185002, 2008.
- [145] T. Pütterich, R. Neu, R. Dux, A. D. Whiteford, M. G. O'Mullane, and the ASDEX Upgrade Team *Plasma Phys. Control. Fusion*, vol. 50, p. 085016, 2008.
- [146] R. Budny, M. Bell, H. Biglari, M. Bitter, C. Bush, C. Cheng, E. Fredrickson, B. Grek, K. Hill, H. Hsuan, A. Janos, D. Jassby, D. Johnson, L. Johnson, B. LeBlanc, D. McCune, D. Mikkelsen, H. Park, A. Ramsey, S. Sabbagh, S. Scott, J. Schivell, J. Strachan, B. Stratton, E. Synakowski, G. Taylor, M. Zarnstorff, and S. Zweben *Nucl. Fusion*, vol. 32, p. 429, 1992.
- [147] M. Oberparleiter, H. Nordman, G. Verdoolaege, and F. Jenko *J. Phys.: Conf. Ser.*, vol. 775, p. 012009, 2016.
- [148] F. W. Perkins in *Heating in Toroidal Plasmas (Proceedings of the 4th International Symposium in Rome)*, vol. 2, p. 977, CEC, 1984.
- [149] P. Ricci, C. Theiler, A. Fasoli, I. Furno, K. Gustafson, D. Irajii, and J. Loizu *Phys. Plasmas*, vol. 18, p. 032109, 2011.
- [150] M. Häse, M. Hirsch, and H. J. Hartfuss *Rev. Sci. Instrum.*, vol. 70, p. 1014, 1999.
- [151] A. E. White, W. A. Peebles, T. L. Rhodes, C. H. Holland, G. Wang, L. Schmitz, T. A. Carter, J. C. Hillesheim, E. J. Doyle, L. Zeng, G. R. McKee, G. M. Staebler, R. E. Waltz, J. C. DeBoo, C. C. Petty, and K. H. Burrell *Phys. Plasmas*, vol. 17, p. 056103, 2010.
- [152] J. Hillesheim, W. Peebles, T. Rhodes, L. Schmitz, A. E. White, and T. Carter *Rev. Sci. Instrum.*, vol. 81, p. 10D907, 2010.
- [153] P. Hennequin, C. Honoré, A. Truc, A. Quemeneur, C. Fenzi-Bonizec, C. Bourdelle, X. Garbet, G. Hoang, and the Tore Supra Team *Nucl. Fusion*, vol. 46, p. S771, 2006.

- [154] T. Happel, T. Görler, P. Hennequin, C. Lechte, M. Bernert, G. D. Conway, S. J. Freethy, C. Honoré, J. R. Pinzon, and U. Stroth *Plasma Phys. Control. Fusion*, vol. 59, p. 054009, 2017.
- [155] S. S. Denk, R. Fischer, O. Maj, E. Poli, J. K. Stober, U. Stroth, B. Vanovac, W. Suttrop, M. Willensdorfer, and the ASDEX Upgrade Team *EPJ Web of Conferences*, vol. 147, p. 02002, 2017.
- [156] P. Rodriguez-Fernandez, A. E. White, A. J. Creely, M. J. Greenwald, N. T. Howard, F. Sciortino, and J. C. Wright *Fusion Sci. Technol.*, vol. 74, p. 65, 2018.
- [157] O. Meneghini, S. P. Smith, L. L. Lao, O. Izacard, Q. Ren, J. M. Park, J. Candy, Z. Wang, C. J. Luna, V. A. Izzo, B. A. Grierson, P. B. Snyder, C. Holland, J. Penna, G. Lu, P. Raum, A. McCubbin, D. M. Orlov, E. A. Belli, N. M. Ferraro, R. Prater, T. H. Osborne, A. D. Turnbull, and G. M. Staebler *Nucl. Fusion*, vol. 55, p. 083008, 2015.
- [158] J. Candy, C. Holland, R. E. Waltz, M. R. Fahey, and E. Belli *Phys. Plasmas*, vol. 16, p. 060704, 2009.
- [159] E. Belli and J. Candy *Plasma Phys. Control. Fusion*, vol. 50, p. 095010, 2008.
- [160] P. B. Snyder, R. J. Groebner, A. W. Leonard, T. H. Osborne, and H. R. Wilson *Phys. Plasmas*, vol. 16, p. 056118, 2009.
- [161] G. L. Jahns, S. K. Wong, R. Prater, S. H. Lin, and S. Ejima *Nucl. Fusion*, vol. 26, p. 226, 1986.
- [162] M. W. Kissick, E. D. Fredrickson, J. D. Callen, C. E. Bush, Z. Chang, P. C. Efthimion, R. A. Hulse, D. K. Mansfield, H. K. Park, J. F. Schivell, S. D. Scott, E. J. Synakowski, G. Taylor, and M. C. Zarnstorff *Nucl. Fusion*, vol. 34, p. 349, 1994.
- [163] M. Soler and J. D. Callen *Nucl. Fusion*, vol. 19, p. 703, 1979.
- [164] J. D. Bell, J. L. Dunlap, V. K. Pare, J. D. Callen, H. C. Howe, E. A. Lazarus, M. Murakami, and C. E. Thomas *Nucl. Fusion*, vol. 24, p. 997, 1984.
- [165] G. L. Jahns, S. K. Wong, R. Prater, S. H. Lin, and S. Ejima, “Measurement of thermal transport by synchronous detection of modulated electron cyclotron heating in the doublet iii tokamak,” Tech. Rep. GA-A17858, GA Technologies, 1985.
- [166] E. D. Fredrickson, J. D. Callen, K. McGuire, J. D. Bell, R. J. Colchin, P. C. Efthimion, K. W. Hill, R. Izzo, D. R. Mikkelsen, D. A. Monticello, V. Pare, G. Taylor, and M. Zarnstorff *Nucl. Fusion*, vol. 26, p. 849, 1986.

- [167] E. D. Fredrickson, K. McGuire, A. Cavallo, R. Budny, A. Janos, D. Monticello, Y. Nagayama, W. Park, G. Taylor, and M. C. Zarnstorff *Phys. Rev. Lett.*, vol. 65, p. 2869, 1990.
- [168] E. D. Fredrickson, A. C. Janos, K. M. McGuire, S. D. Scott, G. Taylor, and Z. Chang *Nucl. Fusion*, vol. 33, p. 1759, 1993.
- [169] F. D. Luca, P. Galli, G. Gorini, A. Jacchia, P. Mantica, N. Deliyannis, M. Erba, and L. Porte *Nucl. Fusion*, vol. 36, p. 909, 1996.
- [170] V. V. Parail, A. Cherubini, J. G. Cordey, M. Erba, P. Galli, R. Giannella, L. Porte, F. Romanelli, A. Rookes, E. M. Springmann, and A. Taroni *Nucl. Fusion*, vol. 37, p. 481, 1997.
- [171] K. W. Gentle, M. E. Austin, and P. E. Phillips *Plasma Sci. Technol.*, vol. 6, p. 2130, 2004.
- [172] P. Mantica and F. Ryter *C. R. Physique*, vol. 7, p. 634, 2006.
- [173] T. Munsat, P. C. Efthimion, B. Jones, R. Kaita, R. Majeski, D. Stutman, and G. Taylor *Phys. Plasmas*, vol. 9, p. 480, 2001.
- [174] Q. Du, B. Ling, A. Ti, X. Gao, Q. Fei, S. Sajjad, and G. Xiong *Plasma Sci. Technol.*, vol. 10, p. 681, 2008.
- [175] H. K. Park, N. C. Luhmann Jr., A. J. H. Donne, I. G. J. Classen, C. W. Domier, E. Mazzucato, T. Munsat, M. J. van de Pol, Z. Xia, and the TEXTOR Team *Phys. Rev. Lett.*, vol. 96, p. 195003, 2006.
- [176] H. K. Park, A. J. H. Donne, N. C. Luhmann Jr., I. G. J. Classen, C. W. Domier, E. Mazzucato, T. Munsat, M. J. van de Pol, Z. Xia, and the TEXTOR Team *Phys. Rev. Lett.*, vol. 96, p. 195004, 2006.
- [177] M. Yamada, Y. Nagayama, W. Davis, E. Fredrickson, A. Janos, and F. Levinton *Rev. Sci. Instrum.*, vol. 63, p. 4623, 1992.
- [178] I. T. Chapman *Plasma Phys. Control. Fusion*, vol. 53, p. 013001, 2011.
- [179] M. E. Tiuri *IEEE Transactions on Military Electronics*, vol. 8, p. 264, 1964.
- [180] W. A. Gardner *Signal Processing*, vol. 29, p. 113, 1992.
- [181] A. J. Creely, K. Ida, M. Yoshinuma, T. Tokuzawa, T. Tsujimura, T. Akiyama, R. Sakamoto, M. Emoto, K. Tanaka, and C. A. Michael *Rev. Sci. Instrum.*, vol. 88, p. 073509, 2017.
- [182] W. Mandl, R. C. Wolf, M. G. von Hellermann, and H. P. Summers *Plasma Phys. Control. Fusion*, vol. 35, p. 1373, 1993.
- [183] H. Weisen *Infrared Phys.*, vol. 25 (3), p. 543, 1985.

- [184] M. Hirsch, E. Holzhauser, J. Baldzuhn, B. Kurzan, and B. Scott *Plasma Phys. Control. Fusion*, vol. 43, p. 1641, 2001.
- [185] G. D. Conway, J. Schirmer, S. Klenge, W. Suttrop, E. Holzhauser, and the ASDEX Upgrade Team *Plasma Phys. Control. Fusion*, vol. 46, p. 951, 2004.
- [186] T. Tokuzawa, A. Ejiri, K. Kawahata, K. Tanaka, I. Yamada, M. Yoshinuma, K. Ida, and C. Suzuki *Rev. Sci. Instrum.*, vol. 83, p. 10E322, 2012.
- [187] G. R. Hanson, J. H. Harris, J. B. Wilgen, C. E. Thomas, S. C. Aceto, L. R. Baylor, J. D. Bell, B. Branas, J. L. Dunlap, and A. C. England *Nucl. Fusion*, vol. 32, p. 1593, 1992.
- [188] R. Sabot, F. Clairet, G. D. Conway, L. Cupido, X. Garbet, G. Falchetto, T. Gerbaud, S. Hacquin, P. Hennequin, S. Heurieux, C. Honoré, G. Leclert, L. Meneses, A. Sirinelli, L. Vermare, and A. Truc *Plasma Phys. Control. Fusion*, vol. 48, p. B421, 2006.
- [189] G. R. McKee, R. J. Fonck, D. K. Gupta, D. J. Schlossberg, M. W. Shafer, R. L. Boivin, and W. Solomon *Plasma Fusion Res.*, vol. 2, p. S1025, 2007.
- [190] Z. Yan, G. R. McKee, R. J. Groebner, P. B. Snyder, T. H. Osborne, and K. H. Burrell *Phys. Rev. Lett.*, vol. 107, p. 055004, 2011.
- [191] R. J. Groebner, C. S. Chang, J. W. Hughes, R. Maingi, P. B. Snyder, X. Q. Xu, J. A. Boedo, D. P. Boyle, J. D. Callen, J. M. Canik, I. Cziegler, E. M. Davis, A. Diallo, P. H. Diamond, J. D. Elder, D. P. Eldon, D. R. Ernst, D. P. Fulton, M. Landreman, A. W. Leonard, J. D. Lore, T. H. Osborne, A. Y. Pankin, S. E. Parker, T. L. Rhodes, S. P. Smith, A. C. Sontag, W. M. Stacey, J. Walk, W. Wan, E. H. J. Wang, J. G. Watkins, A. E. White, D. G. Whyte, Z. Yan, E. A. Belli, B. D. Bray, J. Candy, R. M. Churchill, T. M. Deterly, E. J. Doyle, M. E. Fenstermacher, N. M. Ferraro, A. E. Hubbard, I. Joseph, J. E. Kinsey, B. LaBombard, C. J. Lasnier, Z. Lin, B. L. Lipschultz, C. Liu, Y. Ma, G. R. McKee, D. M. Ponce, J. C. Rost, L. Schmitz, G. M. Staebler, L. E. Sugiyama, J. L. Terry, M. V. Umansky, R. E. Waltz, S. M. Wolfe, L. Zeng, and S. J. Zweben *Nucl. Fusion*, vol. 53, p. 093024, 2013.
- [192] U. Stroth, A. B. Navarro, G. D. Conway, T. Görler, T. Happel, P. Hennequin, C. Lechte, P. Manz, P. Simon, A. Biancalani, E. Blanco, C. Bottereau, F. Clairet, S. Coda, T. Eibert, T. Estrada, A. Fasoli, L. Guimarais, Ö. Gürçan, Z. Huang, F. Jenko, W. Kasperek, C. Koenen, A. Krämer-Flecken, M. E. Manso, A. Medvedeva, D. Molina, V. Nikolaeva, B. Plaum, L. Porte, D. Prisi-azhniuk, T. Ribeiro, B. D. Scott, U. Siart, A. Storelli, L. Vermare, S. Wolf, and the ASDEX Upgrade Team *Nucl. Fusion*, vol. 55, p. 083027, 2015.
- [193] M. Kotschenreuther, W. Dorland, M. A. Beer, and G. W. Hammett *Phys. Plasmas*, vol. 2, p. 2381, 1995.

- [194] T. Uckan, C. Hidalgo, J. D. Bell, J. H. Harris, J. L. Dunlap, J. B. Wilgen, C. P. Ritz, T. L. Rhodes, and A. J. Wootton *Phys. Fluids B*, vol. 3, p. 1000, 1991.
- [195] M. Saffman, S. Zoletnik, N. P. Basse, W. Svendsen, G. Kocsis, and M. Endler *Rev. Sci. Instrum.*, vol. 72 (6), p. 2579, 2001.
- [196] N. P. Basse, P. K. Michelsen, S. Zoletnik, M. Saffman, M. Endler, and M. Hirsch *Plasma Sources Sci. Technol.*, vol. 11, p. A138, 2002.
- [197] T. Happel. Ph.D. thesis, Universidad Carlos III de Madrid, 2010.
- [198] M. Fujiwara, K. Kawahata, N. Ohyabu, O. Kaneko, A. Komori, H. Yamada, N. Ashikawa, L. R. Baylor, S. K. Combs, P. C. deVries, M. Emoto, A. Ejiri, P. W. Fisher, H. Funaba, M. Goto, D. Hartmann, K. Ida, H. Idei, S. Iio, K. Ikeda, S. Inagaki, N. Inoue, M. Isobe, S. Kado, K. Khlopenkov, T. Kobuchia, A. V. Krasilnikov, S. Kubo, R. Kumazawa, F. Leuterer, Y. Liang, J. F. Lyon, S. Masuzaki, T. Minami, J. Miyajima, T. Morisaki, S. Morita, S. Murakami, S. Muto, T. Mutoh, Y. Nagayama, N. Nakajima, Y. Nakamura, H. Nakanishi, K. Narihara, K. Nishimura, N. Noda, T. Notake, S. Ohdachi, Y. Oka, S. Okajima, M. Okamoto, M. Osakabe, T. Ozaki, R. O. Pavlichenko, B. J. Peterson, A. Sagara, K. Saito, S. Sakakibara, R. Sakamoto, H. Sanuki, H. Sasao, M. Sasao, K. Sato, M. Sato, T. Seki, T. Shimosuma, M. Shoji, H. Sugama, H. Suzuki, M. Takechi, Y. Takeiri, N. Tamura, K. Tanaka, K. Toi, T. Tokuzawa, Y. Torii, K. Tsumori, K. Y. Watanabe, T. Watanabe, T. Watari, I. Yamada, S. Yamaguchi, S. Yamamoto, M. Yokoyama, N. Yoshida, Y. Yoshimura, Y. P. Zhao, R. Akiyama, K. Haba, M. Iima, J. Kodaira, T. Takita, T. Tsuzuki, K. Yamauchi, H. Yonezu, H. Chikaraishi, S. Hamaguchi, S. Imagawa, N. Inoue, A. Iwamoto, S. Kitagawa, Y. Kubota, R. Maekawa, T. Mito, K. Murai, A. Nishimura, H. Chikaraishi, K. Takahata, H. Tamura, S. Yamada, N. Yanagi, K. Itoh, K. Matsuoka, K. Ohkubo, I. Ohtake, S. Satoh, T. Satow, S. Sudo, S. Tanahashi, K. Yamazaki, Y. Hamada, and O. Motojima *Nucl. Fusion*, vol. 41 (10), p. 1355, 2001.
- [199] O. Motojima, N. Ohyabu, A. Komori, O. Kaneko, H. Yamada, K. Kawahata, Y. Nakamura, K. Ida, T. Akiyama, N. Ashikawa, W. A. Cooper, A. Ejiri, M. Emoto, N. Ezumi, H. Funaba, A. Fukuyama, P. Goncharov, M. Goto, H. Idei, K. Ikeda, S. Inagaki, M. Isobe, S. Kado, H. Kawazome, K. Khlopenkov, T. Kobuchi, K. Kondo, A. Kostrioukov, S. Kubo, R. Kumazawa, Y. Liang, J. F. Lyon, A. Mase, S. Masuzaki, T. Minami, J. Miyazawa, T. Morisaki, S. Morita, S. Murakami, S. Muto, T. Mutoh, K. Nagaoka, Y. Nagayama, N. Nakajima, K. Nakamura, H. Nakanishi, K. Narihara, Y. Narushima, K. Nishimura, N. Nishino, N. Noda, T. Notake, H. Nozato, S. Ohdachi, Y. Oka, H. Okada, S. Okamura, M. Osakabe, T. Ozaki, B. J. Peterson, A. Sagara, T. Saida, K. Saito, S. Sakakibara, M. Sakamoto, R. Sakamoto, M. Sasao, K. Sato, M. Sato, T. Seki, T. Shimosuma, M. Shoji, H. Suzuki, Y. Takeiri, N. Takeuchi, N. Tamura, K. Tanaka, M. Y. Tanaka, Y. Teramachi, K. Toi,

- T. Tokuzawa, Y. Tomota, Y. Torii, K. Tsumori, K. Y. Watanabe, T. Watari, Y. Xu, I. Yamada, S. Yamamoto, T. Yamamoto, M. Yokoyama, S. Yoshimura, Y. Yoshimura, M. Yoshinuma, N. Asakura, T. Fujita, T. Fukuda, T. Hatae, S. Higashijima, A. Isayama, Y. Kamada, H. Kubo, Y. Kusama, Y. Miura, T. Nakano, H. Ninomiya, T. Oikawa, N. Oyama, Y. Sakamoto, K. Shinohara, T. Suzuki, H. Takenaga, K. Ushigusa, T. Hino, M. Ichimura, Y. Takase, F. Sano, H. Zushi, T. Satow, S. Imagawa, T. Mito, I. Ohtake, T. Uda, K. Itoh, K. Ohkubo, S. Sudo, K. Yamazaki, K. Matsuoka, Y. Hamada, and M. Fujiwara *Nucl. Fusion*, vol. 43, p. 1674, 2003.
- [200] K. Tanaka, L. N. Vyacheslavov, T. Akiyama, A. Sanin, K. Kawahata, T. Tokuzawa, Y. Ito, S. Tsuji-Iio, and S. Okajima *Rev. Sci. Instrum.*, vol. 74 (3), p. 1633, 2003.
- [201] K. Tanaka, C. A. Michael, L. N. Vyacheslavov, A. L. Sanin, K. Kawahata, T. Akiyama, T. Tokuzawa, and S. Okajima *Rev. Sci. Instrum.*, vol. 79, p. 10E702, 2008.
- [202] C. A. Michael, K. Tanaka, L. N. Vyacheslavov, A. Sanin, N. Karchev, T. Akiyama, K. Kawahata, and S. Okajima *Plasma Fusion Res.*, vol. 3, p. S1071, 2008.
- [203] C. A. Michael, K. Tanaka, L. N. Vyacheslavov, A. Sanin, and K. Kawahata *Rev. Sci. Instrum.*, vol. 86, p. 093503, 2015.
- [204] T. Tokuzawa, A. Ejiri, and K. Kawahata *Rev. Sci. Instrum.*, vol. 81, p. 10D906, 2010.
- [205] T. Tokuzawa, S. Inagaki, A. Ejiri, R. Soga, I. Yamada, S. Kubo, M. Yoshinuma, K. Ida, C. Suzuki, K. Tanaka, T. Akiyama, N. Kasuya, K. Itoh, K. Watanabe, H. Yamada, K. Kawahata, and the LHD Experiment Group *Plasma Fusion Res.*, vol. 9, p. 1402149, 2014.
- [206] S. Kubo, H. Idei, T. Shimozuma, Y. Yoshimura, T. Notake, K. Ohkubo, S. Inagaki, Y. Nagayama, K. Narihara, I. Yamada, S. Muto, and S. Morita *AIP Conf. Proc.*, vol. 669, p. 187, 2003.
- [207] T. I. Tsujimura, S. Kubo, H. Takahashi, R. Makino, R. Seki, Y. Yoshimura, H. Igami, T. Shimozuma, K. Ida, C. Suzuki, M. Emoto, M. Yokoyama, T. Kobayashi, C. Moon, K. Nagaoka, M. Osakabe, S. Kobayashi, S. Ito, Y. Mizuno, K. Okada, A. Ejiri, T. Mutoh, and the LHD Experiment Group *Nucl. Fusion*, vol. 55, p. 123019, 2015.
- [208] I. Yamada, K. Narihara, H. Funaba, R. Yasuhara, T. Kohmoto, H. Hayashi, T. Hatae, H. Tojo, T. Sakuma, H. Yoshida, H. Fujita, and M. Nakatsuka *JINST*, vol. 7, p. C05007, 2012.

- [209] C. Suzuki, K. Ida, Y. Suzuki, M. Yoshida, M. Emoto, and M. Yokoyama *Plasma Phys. Control. Fusion*, vol. 55, p. 014016, 2013.
- [210] K. Kawahata, A. Ejiri, K. Tanaka, Y. Ito, and S. Okajima *Fusion Eng. Des.*, vol. 34, p. 393, 1997.
- [211] K. Ida, S. Kado, and Y. Liang *Rev. Sci. Instrum.*, vol. 71 (6), p. 2360, 2000.
- [212] M. Yoshinuma, K. Ida, M. Yokoyama, M. Osakabe, and K. Nagaoka *Fusion Sci. Technol.*, vol. 58, p. 375, 2010.
- [213] S. C. Jardin, N. Pomphrey, and J. DeLucia *Journal of Computational Physics*, vol. 66, p. 481, 1986.
- [214] B. Coppi and N. Sharky *Nucl. Fusion*, vol. 21, p. 1363, 1980.
- [215] W. M. Tang *Nucl. Fusion*, vol. 26, p. 1605, 1986.
- [216] M. Greenwald, D. G. Whyte, P. Bonoli, Z. Hartwig, J. Irby, B. LaBombard, E. Marmor, J. Minervini, M. Takayasu, J. Terry, R. Vieira, A. E. White, S. Wukitch, D. Brunner, R. Mumgaard, and B. Sorbom, “The high-field path to practical fusion energy,” Tech. Rep. RR-18-2, Plasma Science and Fusion Center, 2018.
- [217] 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 *Fusion Engineering and Design*, vol. 100, p. 920, 2015.
- [218] A. Q. Kuang, N. M. Cao, A. J. Creely, C. A. Dennett, J. Hecla, B. LaBombard, R. A. Tinguely, E. A. Tolman, H. Hoffman, M. Major, J. Ruiz Ruiz, D. Brunner, P. Grover, C. Laughman, B. N. Sorbom, and D. G. Whyte *Fusion Engineering and Design*, vol. 137, p. 221, 2018.
- [219] R. T. Mumgaard, M. Greenwald, J. P. Freidberg, S. M. Wolfe, Z. S. Hartwig, D. Brunner, B. N. Sorbom, and D. G. Whyte, “Scoping study for compact high-field superconducting net energy tokamaks,” (San Jose, California), American Physical Society Division of Plasma Physics, 2016.
- [220] J. Wesson, *Tokamaks*. International series of monographs on physics, Oxford University Press, 4th ed., 2011.
- [221] F. Troyon, R. Gruber, H. Saurenmann, S. Semenzato, and S. Succi *Plasma Phys. Control. Fusion*, vol. 26, p. 209, 1984.
- [222] A. R. Polevoi, V. S. Mukhovatov, M. Shimada, S. Y. Medvedev, A. A. Ivanov, A. S. Kukushkin, and Y. Murakami in *Proceedings of joint conference of ITC-12 and APFA '01: Frontiers in plasma confinement and related engineering/plasma science*, 2003.

- [223] A. J. Creely, L. M. Milanese, E. A. Tolman, J. H. Irby, S. B. Ballinger, S. Frank, A. Q. Kuang, B. L. Linehan, W. McCarthy, K. J. Montes, T. Mouratidis, J. F. Picard, P. Rodriguez-Fernandez, A. M. Rosenthal, A. J. Sandberg, F. Sciortino, R. A. Simpson, R. A. Tinguely, M. Zhou, and A. E. White *Fusion Engineering and Design*, vol. Submitted, p. NA, 2019.
- [224] V. V. Mirnov, W. X. Ding, D. L. Brower, M. A. Van Zeeland, and T. N. Carlstrom *Phys. Plasmas*, vol. 14, p. 102105, 2007.
- [225] H. Soltwisch *Rev. Sci. Instrum.*, vol. 57, p. 1939, 1986.
- [226] W. F. Bergerson, P. Xu, J. H. Irby, D. L. Brower, W. X. Ding, and E. S. Marmor *Rev. Sci. Instrum.*, vol. 83, p. 10E316, 2012.
- [227] C. P. Kasten, J. H. Irby, R. Murray, A. E. White, and D. C. Pace *Rev. Sci. Instrum.*, vol. 83, p. 10E301, 2012.
- [228] M. A. Van Zeeland, T. N. Carlstrom, D. K. Finkenthal, R. L. Boivin, A. Colio, D. Du, A. Gattuso, F. Glass, C. M. Muscatello, R. O'Neill, M. Smiley, J. Vasquez, M. Watkins, D. L. Brower, J. Chen, W. X. Ding, D. Johnson, P. Mauzey, M. Perry, C. Watts, and R. Wood *Plasma Phys. Control. Fusion*, vol. 59, p. 125005, 2017.
- [229] H. Q. Liu, J. P. Qian, Y. X. Jie, W. X. Ding, D. L. Brower, Z. Y. Zou, W. M. Li, H. Lian, S. X. Wang, Y. Yang, L. Zeng, T. Lan, Y. Yao, L. Q. Hu, X. D. Zhang, and B. N. Wan *Rev. Sci. Instrum.*, vol. 87, p. 11D903, 2016.
- [230] A. J. H. Donné, M. F. Graswinckel, M. Cavinato, L. Giudicotti, E. Zilli, C. Gil, H. R. Koslowski, P. McCarthy, C. Nyhan, S. Prunty, M. Spillane, and C. Walker *Rev. Sci. Instrum.*, vol. 75, p. 4694, 2004.
- [231] A. Mlynek, L. Casali, O. Ford, H. Eixenberger, and the ASDEX Upgrade Team *Rev. Sci. Instrum.*, vol. 85, p. 11D408, 2014.
- [232] G. Braithwaite, N. Gottardi, G. Magyar, J. R. J. O'Rourke, and D. Veron *Rev. Sci. Instrum.*, vol. 60, p. 2825, 1989.
- [233] A. J. H. Donné, A. E. Costley, R. Barnsley, H. Bindslev, R. Boivin, G. Conway, R. Fisher, R. Giannella, H. Hartfuss, M. G. von Hellermann, E. Hodgson, L. C. Ingesson, K. Itami, D. Johnson, Y. Kawano, T. Kondoh, A. Krasilnikov, Y. Kusama, A. Litnovsky, P. Lotte, P. Nielsen, T. Nishitani, F. Orsitto, B. J. Peterson, G. Razdobarin, J. Sanchez, M. Sasao, T. Sugie, G. Vayakis, V. Voit-senya, K. Vukolov, C. Walker, K. Young, and the ITPA Topical Group on Diagnostics *Nucl. Fusion*, vol. 47, p. S337, 2007.
- [234] K. Kawahata, T. Akiyama, K. Tanaka, S. Okajima, and K. Nakayama *Journal Phys.: Conf. Ser.*, vol. 227, p. 012022, 2010.

- [235] K. Kawahata, T. Akiyama, K. Tanaka, K. Nakayama, and S. Okajima *JINST*, vol. 7, p. C02002, 2012.
- [236] J. Chen, G. Zhuang, Q. Li, Y. Liu, L. Gao, Y. N. Zhou, X. Jian, C. Y. Xiong, Z. J. Wang, D. L. Brower, and W. X. Ding *Rev. Sci. Instrum.*, vol. 85, p. 11D303, 2014.
- [237] M. A. Van Zeeland, R. L. Boivin, T. N. Carlstrom, and T. M. Deterly *Rev. Sci. Instrum.*, vol. 79, p. 10E719, 2008.
- [238] M. A. Van Zeeland, R. L. Boivin, D. L. Brower, T. N. Carlstrom, J. A. Chavez, W. X. Ding, R. Feder, D. Johnson, L. Lin, R. C. O'Neill, and C. Watts *Rev. Sci. Instrum.*, vol. 84, p. 043501, 2013.
- [239] C. P. Kasten, A. E. White, and J. H. Irby *Phys. Plasmas*, vol. 21, p. 042305, 2014.
- [240] W. W. Heidbrink *Phys. Plasmas*, vol. 15, p. 055501, 2008.