

Impact of Wavelength Scale Density Variation on Microwave Propagation in Tokamak Plasmas

Lou Holland

Doctor of Philosophy

University of York

School of Physics, Engineering and Technology

September 2023

Abstract

The propagation of microwaves through magnetised plasmas in the presence of wavelength scale density variation poses an interesting physics problem. Microwaves have many uses in tokamaks, from diagnostics that help characterise the plasma to high-power beams used for heating and current drive. It is therefore important to be able to accurately predict the path that they will take in tokamak plasmas, even in the presence of fluctuations.

To this end, a full-wave cold-plasma code utilising the FDTD method (EMIT) has been developed both in 2D and 3D. EMIT-2D minimises computational cost, allowing full simulation of the beam from the antenna to the absorption region. A benchmark of the code was carried out before it was applied to the problem of OX-mode conversion. In plasmas with steep density gradients, mode conversion efficiency was found to decrease sharply due to the converted X-mode tunnelling back out of the plasma.

EMIT-2D was also used in a study of ECRH beam broadening by turbulence on DIII-D. Significant beam broadening was measured experimentally in three operating scenarios. Diagnostic data was used to generate synthetic turbulent density profiles for simulations. The simulations agreed with experiment, providing a direct comparison between simulation and experimental measurements of beam broadening for the first time, but diagnostic uncertainty led to significant uncertainty in the simulated results, motivating the need for future turbulence diagnostics of better spatial resolution.

To further characterise how beam broadening by electrostatic turbulence depends on plasma and beam parameters, a series of parameter scans were carried out covering tokamak relevant parameter ranges. The parameter scans were conducted in pairwise combinations of the parameters in order to determine the separability of the dependencies, and an empirical formula was found for fusion-relevant scenarios allowing the prediction of beam-broadening in microseconds instead of the hours required for full-wave simulations.

Contents

Abstract	2
List of Tables	7
List of Figures	10
Acknowledgements	20
Declaration	21
1 Introduction	23
1.1 Magnetically Confined Fusion	23
1.1.1 Motivation for Fusion	23
1.1.2 The Tokamak	24
1.1.3 The Spherical Tokamak	25
1.2 Turbulence in Fusion Plasmas	26
1.3 Microwaves in Fusion Plasmas	28
1.3.1 Heating and current drive	28
1.3.2 Diagnostics	30
1.4 Thesis Outline	31
2 Waves in Plasmas	33
2.1 Waves in Cold Plasmas	34
2.1.1 Propagation Parallel to Magnetic Field	39
2.1.2 Propagation Perpendicular to Magnetic Field	41
2.2 Limitations of the Cold Model	44

3	Numerical Methods	45
3.1	Summary of Approaches	45
3.1.1	Ray Tracing	45
3.1.2	Full-Wave Modelling	46
3.1.3	Kinetic Modelling	47
3.2	The FDTD Method for Full-Wave Modelling in EMIT-3D	48
3.2.1	The Yee Algorithm	48
3.2.2	The Plasma Response	51
3.2.3	The Antenna	53
3.2.4	Boundary Conditions	55
4	Code Development: EMIT	57
4.1	EMIT-3D	57
4.1.1	Change to Update Equations	57
4.1.2	Elliptical polarisation	58
4.2	EMIT-2D	59
4.2.1	Update Equations in 2D	60
4.2.2	Antenna in 2D	61
4.2.3	Parallelisation in OpenMP	62
4.2.4	Convergence Testing, Choice of Timestep and Resolution, and Numerical Error	62
4.2.5	Benchmarking	64
5	Code Comparison and OX Mode Conversion Study	67
5.1	Code Comparison	68
5.1.1	Simulation Set-Up	68
5.1.2	Scenario 1	70
5.1.3	Scenario 2	71
5.1.4	Scenario 3	75
5.1.5	Scenario 4	75
5.1.6	Scenario 5	79
5.1.7	Scenario 6	82
5.2	Effect of Density Scale Length on OX Mode Conversion	84

<i>CONTENTS</i>	5
5.3 Summary and Conclusions	87
6 ECRH Beam Broadening on DIII-D	89
6.1 Motivation	89
6.2 Experimental Measurements	91
6.3 Simulations	93
6.3.1 Fluid vs Synthetic Turbulence	94
6.3.2 Effect of Core Turbulence	98
6.3.3 Final Simulation Set-Up	99
6.4 Results	100
6.5 Conclusions and Further Work	102
7 Parametric Dependence of Beam Broadening by Turbulence	105
7.1 Introduction	105
7.2 Simulation Set Up	108
7.2.1 Turbulent Density Profiles	108
7.2.2 Calculating Beam Broadening	111
7.2.3 Parameter Scans	112
7.2.4 Numerical Parameters	113
7.2.5 Fitting	113
7.3 Parametric Dependence	115
7.4 Conclusions and Further Work	125
8 Conclusions and Future Work	129
8.1 Conclusions	130
8.1.1 Code Development	130
8.1.2 Benchmark and OX Mode Conversion Study	131
8.1.3 ECRH Beam Broadening on DIII-D	131
8.1.4 Parametric Dependence of Beam Broadening by Turbulence	132
8.2 Future Work	134
8.2.1 Beam Broadening by Turbulence	134
8.2.2 EMIT-3D Code Development	135
A Parametric Dependence Fit Parameters	139

List of References

141

List of Tables

5.1	A summary of the scenarios used to benchmark the codes against each other. The parameters x_{ant} and z_{ant} define the location of the antenna, and θ is the angle of propagation of the beam to the z -axis. All other symbols have their previously given meanings.	68
5.2	A summary of the numerical parameters used in the benchmark scenarios by EMIT-2D. Resolution is the number of Yee cells per vacuum wavelength, and N_x and N_z are the number of Yee cells over which the simulation domain extends. CFL relates the time step to the grid spacing, and N_t is the total number of timesteps over which the simulation was run. Finally, ν is the numerical collision frequency normalised to the vacuum wave frequency. . . .	69
6.1	Configuration of relevant discharges on DIII-D, all with $B_t = 2$ T. The plasma current I_p is in units of MA and the background plasma density at a normalised radius of $\rho = 0.95$ is in units of $\times 10^{19} \text{ m}^{-3}$. Parameters κ and δ describe the shaping of the plasma, and ρ_{ECH} is the normalised radius of the power deposition location. The broadening factor, b , comes from the width of the deposition profile (after transport analysis) divided by the width predicted by TORAY-GA. This broadening can be seen to scale with fluctuation level $\delta n_e/n_e$ at $\rho = 0.95$	93

- 6.2 A summary of the numerical parameters used in the benchmark scenarios. Resolution is the number of Yee cells per vacuum wavelength, and N_x and N_z are the number of Yee cells over which the simulation domain extends. CFL relates the time step to the grid spacing, and N_t is the total number of timesteps over which the simulation was run. Finally, ν is the numerical collision frequency outside of the damping boundaries normalised to the vacuum wave frequency. 100
- 6.3 Comparison of experimentally measured (Exp) and simulated broadening. Uncertainty on simulated results comes from the standard error on the mean from the ensemble of turbulence profiles for each scenario. The ‘Default Sim’ is the simulated broadening when using the density and fluctuation data as measured rather than at the extremal ends of the possible range. Uncertainty in the diagnostic data used as input for the simulations is used to determine the minimum (Min Sim) and maximum (Max Sim) possible simulated broadening. This leads to large uncertainty on the final simulated (Fin Sim) result, calculated as the midpoint of the range of results, but ultimately yields agreement with the experimental results. 102
- 7.1 Parameters that are varied, how they are varied, and the ‘base’ value each parameter takes in scans where it is not varied. Note that the background density is only varied through nine values instead of ten, as at the timescales needed for higher densities, the stability condition proved prohibitively computationally expensive. 112
- 7.2 Number of timesteps that the simulation was run for based on background density, assuming a resolution of 25 Yee cells per wavelength. Higher background densities result in the simulation taking longer to reach steady state. In all cases, the code was run for at least one full wave period after steady state was reached. 114

A.1 Look up table for beam broadening by plasma turbulence based on plasma and beam parameters. p is the parameter (or combination of parameters where the dependence is not separable), f is the value of the fit function, and e is the error on that fit. $p_1 = (L_{\text{perp}}/\lambda_0)^{0.4}/(L_{\text{bin}}/\lambda_0)^{0.9}$, $p_2 = (L_{\text{perp}}/\lambda_0)^{0.9}(L_{\text{bin}}/\lambda_0)^{0.4}$, $p_3 = 0.2(n_{e,0}/n_{\text{crit}}) - 1.0(\delta n_e/n_{e,0})$, $p_4 = 1.0(n_{e,0}/n_{\text{crit}}) + 0.2(\delta n_e/n_{e,0})$, $p_5 = w_0/\lambda_0$, and $p_6 = W_{\text{turb}}/\lambda_0$ 140

List of Figures

1.1	Configuration of the magnetic field in a tokamak. Large D-shaped magnets generate the toroidal field. A plasma current generates the poloidal field, aided by inner and outer poloidal field coils. The resulting magnetic field follows a helical pattern around the torus. Figure courtesy of Eurofusion [1].	25
1.2	Different plasma shapes and volumes for a conventional and spherical tokamak. Figure courtesy of UKAEA [2].	26
3.1	A Yee unit cell showing the grid point locations at which values are calculated. Each value is calculated using its nearest neighbours on the previous time step. The grid locations of \mathbf{E} and $\mathbf{H} = \mathbf{B}/\mu$ are shown. Figure reproduced from [3].	49
4.1	Speedup of EMIT-2D on the MARCONI supercomputer [4] for two different simulation domain sizes shown by the blue and orange data points. The black dashed line represents the ideal speedup where a doubling of cores leads to a halving of wall time. Near perfect speedup can be seen with only slight deviation after 40 cores is exceeded. Running for 5000 timesteps, the runs on the smallest number of cores (one and four cores respectively) took approximately 8 hours of wall time.	63

- 4.2 Comparison of beam broadening results when halving the time-step. The ensemble average shows the average beam after travelling through an ensemble of turbulence profiles. The background shows the beam at the same location in the simulation domain after propagating through an equivalent background plasma with no turbulence present. The broadening is found by comparing the FWHM of the ensemble average and the background beam. The uncertainty values come from the standard error on the mean of the ensemble average. From this comparison, we can see that halving the timestep decreases the broadening by 0.0034, equating to a numerical error or around 0.1 %, which is significantly smaller than the uncertainty on the result. The numerical uncertainty due to discretisation can therefore be neglected as insignificant compared to other sources of uncertainty. 65
- 4.3 Relationship between normalised plasma density and wavelength of (a) an O-mode (left) and (b) an X-mode (right) wave propagating through the plasma. The black dots are from simulations in EMIT-2D, and the red lines are calculated directly from the dispersion relations given by Eqs. (2.53) and (2.54) respectively. 66
- 5.1 Simulation domain for scenarios 4, 5 and 6. The colourmap indicates normalised density $X = n_e/n_{\text{crit}}$. An arrow indicates the direction of the background field, $B_0 = 0.85$ T corresponding to $Y \approx 0.5$, and a star indicates the position of the antenna. 69
- 5.2 Full-wave simulation results for scenario 1. A Gaussian beam is launched with its focus at the antenna, propagating in the z -direction through a vacuum. The power in the beam is read out at different slices in z and compared between the codes. 72
- 5.3 Percentage change in the total power in the electric field at different slices in z for EMIT-2D scenarios 1 and 2. The point corresponding to the z -slice within the damping boundaries is omitted to avoid skewing the axis of the plot. The percentage change outside of the damping boundaries does not exceed a magnitude of 0.31 % of the original value in the antenna plane for any of the scenarios. 73

5.4 Full-wave simulation results for scenario 2. A Gaussian beam is launched with its focus 0.2 m into the plasma, propagating in the z -direction through a vacuum. The power in the beam is read out at different slices in z and compared between the codes. 74

5.5 Full-wave simulation results for scenario 3. A Gaussian beam is launched with its focus at the antenna, propagating at 30° to the z -axis through a vacuum. The power in the beam is read out at different slices in z and compared between the codes 76

5.6 Full-wave simulation results for scenario 4. A Gaussian beam is launched in vacuum with its focus at the antenna, propagating at 30° to the z -axis into a plasma with a linear density gradient. It is reflected at the cut-off density in the plasma before propagating back to the antenna plane. The power in the beam is read out at different slices in z and compared between the codes . . . 77

5.7 Full-wave simulation results for scenario 4 using EMIT-2D with two different resolutions. A Gaussian beam is launched in vacuum with its focus at the antenna, propagating at 30° to the z -axis into a plasma with a linear density gradient. It is reflected at the cut-off density in the plasma before propagating back to the antenna plane. The power in the beam is read out at different slices in z and compared between the resolutions. The top figure shows microwave beam power at various slices through the simulation domain when using a resolution of 50 Yee Cells per wavelength. The middle figure shows the microwave beam power at various slices through the simulation domain when using a resolution of 100 Yee Cells per wavelength. The bottom figure shows the difference between the two at various slices through the simulation domain. This difference never exceeds 1.2% and is highest in the slice closest to the cut-off, where we expect constructive and destructive interference to be occurring, making the code more sensitive to slight changes in relative position between the cut-off surface and the location we are reading the field due to the discrete grid. 78

- 5.8 Full-wave simulation results for scenario 4 using EMIT-2D with three different polarisations. A Gaussian beam is launched in vacuum with its focus at the antenna, propagating at 30° to the z -axis into a plasma with a linear density gradient. It is reflected at the cut-off density in the plasma before propagating back to the antenna plane. The power in the beam is read out at different slices in z 80
- 5.9 Full-wave simulation results for scenario 4, now using an elliptical polarisation in EMIT-2D. A Gaussian beam is launched in vacuum with its focus at the antenna, propagating at 30° to the z -axis into a plasma with a linear density gradient. It is reflected at the cut-off density in the plasma before propagating back to the antenna plane. The power in the beam is read out at different slices in z and compared between the codes. 81
- 5.10 Full-wave simulation results for scenario 5. A diverging O-mode Gaussian beam is launched in vacuum with uniform background magnetic field in the x -direction, propagating at 30° to the z -axis into a plasma with a linear density gradient. At the O-mode cut-off density, the O-mode is reflected but some mode conversion to the X-mode also occurs. The power in the beam is read out at different slices in z and compared between the codes 83
- 5.11 Full-wave simulation results for scenario 6. A diverging O-mode Gaussian beam is launched in vacuum with uniform background magnetic field in the x -direction, propagating at the optimum angle for OX mode conversion efficiency, which equates to an angle of 47° to the z -axis into a plasma with a linear density gradient. At the O-mode cut-off density, mode conversion to the X-mode occurs, and any O-mode that is not converted is reflected back out of the plasma. The power in the beam is read out at different slices in z and compared between the codes 85

- 5.12 Snapshot of the wave power for scenario 6. Colour represents power in arbitrary units on a logarithmic scale. The diverging Gaussian beam can be seen travelling toward the O-mode cut-off at the optimum angle for mode conversion. At the O-mode cut-off, the majority of the beam is converted to the SX-mode. In a warm plasma, this would then travel to the UHR where it converted to an EBW before propagating further into the plasma. but, in this cold plasma code, the converted X-mode is instead dissipated with numerical collisionality. The tails of the injected O-mode beam do not have the optimum wave-vector so are not converted to the SX-mode. Any unconverted O-mode is reflected back, resulting in the hollow middle of the reflected beam. 86
- 5.13 How OX mode conversion efficiency depends on density scale length of a linear density profile in the presence of uniform background magnetic field. 87
- 5.14 Power in the antenna plane for two scenarios with different density scale lengths. The large peak to the left in each plot is the emitting antenna. The smaller peaks to the right are the reflections, indicating the parts of the beam that were not successfully mode converted. The gap in the middle of the reflected beam in figure (a) is typical for a case where OX mode conversion has taken place. This is absent from figure (b) due to the X-mode tunnelling back out of the plasma along with the reflected O-mode. 88
- 6.1 Cross-section of the DIII-D tokamak showing microwave launch. The dashed black lines represent closed flux surfaces, the solid black line represents the separatrix, and the thick grey line indicates the vacuum vessel. The injected microwave beam is illustrated by the blue arrow with the 2nd harmonic electron cyclotron resonance indicated by the green line. The red rectangle shows the spatial domain of the 2D full-wave simulations that were carried out to compare to experiment. 92
- 6.2 E_{RMS} profile of microwave beam at absorption region perpendicular to B_0 and direction of beam propagation. Broadening is calculated by comparing to a background case where no turbulence was present. 95

- 6.3 Background density and fluctuation envelope for the L-mode case, as was used by for the fluid turbulence vs synthetic turbulence comparison. It should be noted that the actual simulation domain started at the outer edge and is cut off at $R = 2.18$ m. 96
- 6.4 Comparing the broadening found running simulations through fluid and synthetic turbulence with the same macroscopic properties of turbulence correlation lengths, background density, and fluctuation level envelope. The top panel shows the normalised density n_e/n_{crit} for a single turbulent snapshot. The middle pannel shows a single snapshot of the wave electric field passing through a single turbulent snapshot. The bottom panel shows the RMS electric field read out at the backplane of the simulation. Axis are plotted in grid points where a resolution of 25 grid points per vacuum wavelength was used. For the fluid turbulence, the broadening is found to be 1.3 ± 0.1 and for the synthetic turbulence, it's found to be 1.2 ± 0.1 97
- 6.5 Simulation set-up for inclusion of core turbulence to the diverted L-mode case. Top panel shows the background density, middle panel shows the normalised fluctuation level envelope with and without the inclusion of core turbulence at a level of 1% of the background density. The bottom data shows the fluctuation amplitude normalised to the critical density 99
- 6.6 Background density and fluctuation level envelopes for all three scenarios, as used to compare to experiments. The experimental data used to create these profiles are shown as dots. 101
- 7.1 Example ensemble average background density profile and fluctuation level envelope (top), an example turbulent density profile (bottom left) and the RMS Electric field (bottom right) of a microwave beam propagating through it. The simulation domain is given in vacuum wavelengths. The colour represents the normalised density $X = \omega_{\text{pe}}^2/\omega^2 = n_{e,0}/n_{\text{crit}}$ (bottom left) and the RMS field strength in arbitrary units (bottom right). Arrows are marked on to indicate the direction of the background field as well as the direction in which the two different turbulence correlation lengths are defined. 109

- 7.2 An example of the RMS electric field at the back-plane of the simulation domain, after travelling through the turbulent layer, averaged over an ensemble of 20 uncorrelated turbulence profiles. This is the "base case" as defined in Table 7.1. A Gaussian fit to the ensemble average is performed and compared to the background Gaussian beam to find the broadening factor. 111
- 7.3 Dependence of broadening on microwave beam waist (w_0) and width of the turbulence layer (W_{turb}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present. 116
- 7.4 Dependence of broadening on turbulence correlation length in the direction perpendicular to the magnetic field (L_{perp}) and in the direction binormal to both the magnetic field and direction of beam propagation (L_{bin}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present. From this, we see that both correlation lengths affect broadening differently, and the tilted nature indicates the dependence is not separable. 116
- 7.5 Dependence of broadening on background density ($n_{e,0}/n_{\text{crit}}$) and fluctuation amplitude ($\delta n_e/n_{e,0}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present. From this, we see that broadening increases with both background density and fluctuation amplitude, and the tilted nature indicates the dependence is not separable. 117
- 7.6 Dependence of broadening on microwave beam waist (w_0) and fluctuation level ($dn_e/n_{e,0}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present. 117
- 7.7 Dependence of broadening on width of turbulence layer (W_{turb}) and fluctuation level ($dn_e/n_{e,0}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present. 118

7.8	Dependence of broadening on binormal correlation length (L_{bin}) and fluctuation level ($dn_e/n_{e,0}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.	118
7.9	Dependence of broadening on binormal correlation length (L_{bin}) and width of turbulence layer (W_{turb}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.	119
7.10	Dependence of broadening on perpendicular correlation length (L_{perp}) and fluctuation level ($dn_e/n_{e,0}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.	119
7.11	Dependence of broadening on perpendicular correlation length (L_{perp}) and width of turbulence layer (W_{turb}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.	120
7.12	Dependence of broadening on binormal correlation length (L_{bin}) and background density ($n_{e,0}/n_{\text{crit}}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.	120
7.13	Dependence of broadening on perpendicular correlation length (L_{perp}) and background density ($n_{e,0}/n_{\text{crit}}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.	121
7.14	Dependence of broadening on width of the turbulence layer (W_{turb}) and background density ($n_{e,0}/n_{\text{crit}}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.	121
7.15	Dependence of broadening on microwave beam waist (w_0) and binormal correlation length (L_{bin}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.	122

- 7.16 Dependence of broadening on microwave beam waist (w_0) and perpendicular correlation length (L_{perp}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present. 122
- 7.17 Dependence of broadening on microwave beam waist (w_0) and background density ($n_{e,0}/n_{\text{crit}}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present. 123
- 7.18 The results for a beam broadening simulation through an ensemble of turbulent profiles with all values as their base case apart from $W_{\text{turb}}/\lambda_0 = 50$. The Ensemble average E_{RMS} can be seen to be very low and flat such that it starts to look similar to background noise rather than a discernible Gaussian beam profile. 125
- 7.19 The normalised dependence of broadening on each parameter considered. For the case where the dependence of two parameters is not separable, the dependence on a linear combination of those parameters is shown instead. The orange points correspond to the base value of each parameter, which was used in scans where other parameters were varied. The fit parameters are normalised to this point. The solid lines serve as a guide to the eye. 126

Acknowledgments

I would like to thank my supervisor Roddy Vann for the opportunity to get involved with such a fascinating area of research along with my collaborator Alf Köhn-Seeman for all the support and knowledge they both shared along the way. In addition, I want to say how grateful I am to David Woodward for the time he took at the beginning of my PhD to help get me up to speed with running the code – one week at the University of Strathclyde with his help probably saved me months of work. And of course, thanks to my collaborators at DIII-D, without whom one of the chapters of this thesis would not have been possible.

I would also like to thank my parents and step-parents, who provided support and encouragement throughout my formative years and without whom I'd never have been in a position to complete a PhD. And I would like to give special thanks to my partner for following me to York and providing endless emotional support throughout my PhD, especially throughout the lockdown period.

This author acknowledges access to the EUROfusion High-Performance Computer (Marconi-Fusion) through EUROfusion. This work also used the ARCHER UK National Supercomputing Service, allocated through the Plasma HEC (EP/R029148/1), the ARCHER2 UK National Supercomputing Service, allocated through TDoTP (EP/R034737/1), the Cirrus UK National Tier-2 HPC Service at EPCC funded by the University of Edinburgh and EPSRC (EP/P020267/1). Simulations were also conducted on the Viking Cluster, which is a high-performance computing facility provided by the University of York. We are grateful for computational support from the University of York High Performance Computing service, Viking and the Research Computing team.

This research was funded by the EPSRC Centre for Doctoral Training in Science and Technology of Fusion Energy grant EP/L01663X.

Declaration

This thesis has not previously been accepted for any degree and is not being concurrently submitted in candidature for any degree other than Doctor of Philosophy of the University of York. This thesis is the result of my own investigations, except where otherwise stated. All other sources are acknowledged by explicit references.

I declare that the work presented in this thesis, except where it is otherwise stated, is based on my own research and has not been submitted previously for a degree in this or any other university. All sources are acknowledged as references. Parts of the work presented in this thesis have been published in:

- A. Köhn-Seeman, B. E. Eliasson, S. J. Freethy, L. A. Holland and R.G.L. Vann, ‘*Benchmarking full-wave codes for studying the O-SX mode conversion in MAST Upgrade*’, EPJ Web of Conferences **277**, 01010 (2023)
- M.W. Brookman, L.A. Holland, M.B. Thomas, M.E. Austin, K. Barada, K.W. Gentle, R.J. La Haye, J.B. Leddy, C.C. Petty, T.L. Rhodes, ‘*Broadening of microwave heating beams in the DIII-D tokamak by edge turbulence*’, Nuclear Fusion **63** 044001 (2023)
- L.A. Holland, A. Köhn-Seeman, R.G.L. Vann, ‘*Parametric dependence of microwave beam broadening by plasma density turbulence*’, Nuclear Fusion **63** 056013 (2023)

The simulation results in Chapter 5 from codes IPF-FDMC and FFW were produced by A. Köhn-Seeman and B. E. Eliasson respectively for the purpose of a benchmark between our three codes.

I also wish to declare that in the final year of study for this PhD, I started part-time work as a Research Trainee in Microwave Current Drive on the STEP project. None of the research I have done in this role has formed part of the thesis.

Chapter 1

Introduction

1.1 Magnetically Confined Fusion

1.1.1 Motivation for Fusion

As demand for energy continues to increase, and climate change drives the need for alternatives to fossil fuels, nuclear fusion remains a promising alternative means of energy production. There are numerous potential advantages to the use of nuclear fusion to generate power. The reaction of interest



does not produce any greenhouse gases so will not contribute to the growing climate crisis, the fuel sources also aren't at pressing risk of depletion, and no high-level radioactive waste is produced [5]. There are other possible fusion reactions, but the one shown in Eq. 1.1 is the most favourable one.

In order to generate electricity from fusion, the reactor needs to generate more energy than it takes to run. One of the major hurdles to overcome is making the reaction self-sustaining, such that all of the energy needed to maintain the required temperature comes from the energetic ${}^4_2\text{He}$ ions that remain within the plasma. Balancing this need against the energy that will be lost via particles escaping confinement and radiative losses produces an inequality known as the Lawson criterion [6]

$$nT\tau_E > 3 \times 10^{21} \text{m}^{-3} \text{keVs} \quad (1.2)$$

which is used to define the ‘ignition point’ at which external heating is no longer required as the plasma is self-sustaining. Here, n is the particle density in m^{-3} , T is the temperature in keV, and τ_E is the energy confinement time in s.

Two main approaches to achieve this are inertial confinement and magnetic confinement. This thesis is focused on magnetic confinement fusion (MCF), where the Deuterium and Tritium fuel is ionised to form a plasma, which is confined by magnetic fields. The aim is then to hold the plasma in steady state for long enough, and at sufficiently high temperatures and densities, that enough fusion reactions occur that there is net energy gain. The three main types of device used for this purpose are tokamaks, spherical tokamaks, and stellarators. Stellarators have a more complicated magnetic field so, while a lot of the fundamental physics is unchanged, pose some different and unique issues when it comes to the study of microwave propagation in stellarator plasmas. The focus from this point on will be on tokamaks and spherical tokamaks which, while posing their own separate and unique issues, also have a number of similarities.

1.1.2 The Tokamak

The tokamak is an MCF device that uses magnetic fields to confine the plasma in a torus (or doughnut) shape. The Deuterium and Tritium fuel is ionised to form a plasma made up of positive nuclei and negative electrons. These charged particles orbit and stream along magnetic field lines so one might suppose that a simple closed loop of field lines (a purely toroidal field) would effectively confine the plasma.

However, the gradient and curvature of the magnetic field introduce drifts in the particle velocity known as grad-B drift and curvature drift. Grad-B drift arises as the radius of a charged particle’s orbit around a field line is dependent on magnetic field strength. A particle orbiting over a region where there’s a gradient in B therefore has a distorted orbit, with a smaller gyro-radius in the region of higher field, and a larger gyro-radius in the region of lower field. This leads to a net drift perpendicular to the magnetic field gradient. Curvature drift arises as, in order for a particle to travel along a curved field line, it requires a centripetal force perpendicular to that field line. As no such force exists in the plasma, a drift velocity away from the field line arises. These drifts mean that a purely toroidal field is not sufficient to confine the plasma effectively. Instead, an additional poloidal field is required. This is achieved by driving a toroidal current in the plasma, which then generates

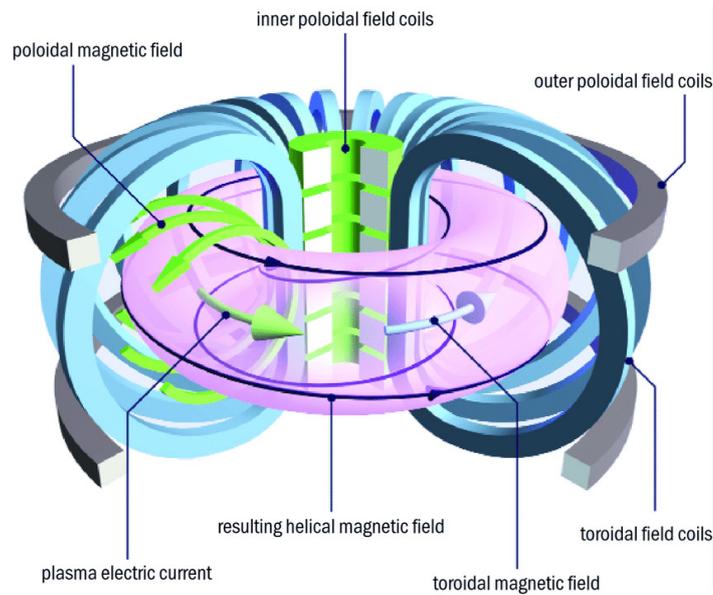


Figure 1.1: Configuration of the magnetic field in a tokamak. Large D-shaped magnets generate the toroidal field. A plasma current generates the poloidal field, aided by inner and outer poloidal field coils. The resulting magnetic field follows a helical pattern around the torus. Figure courtesy of Eurofusion [1].

a poloidal magnetic field. The resultant magnetic field configuration can be seen in Fig. 1.1 [1].

In existing tokamaks, the required plasma current is predominantly generated via ramping up the field in the central solenoid of the device. However, this will not be sufficient for a steady-state device as the field cannot be ramped up indefinitely, limiting the length of operation to the duration of the ramp up. Therefore, additional methods of plasma current generation will be required to drive and maintain the current during steady-state operation. Potential current drive systems include neutral beam injection, and the use of radio frequency waves, as will be discussed in more detail in Sec. 1.3.1.

1.1.3 The Spherical Tokamak

Spherical tokamaks are similar to conventional tokamaks in their overall rough shape but have a smaller aspect ratio $A = R/a$, where R is major radius and a is minor radius of the device. The result is a plasma configuration that looks less like the doughnut in a conventional tokamak and more like a cored apple. The difference between the plasma shapes and volumes of the two devices can be seen in Fig. 1.2.

These devices are significantly smaller than conventional tokamaks while still being ca-

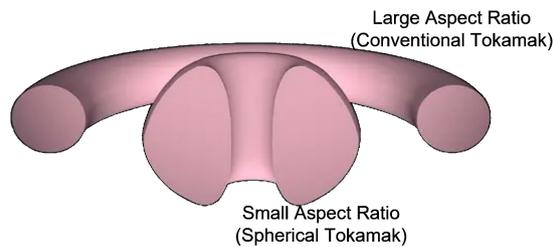


Figure 1.2: Different plasma shapes and volumes for a conventional and spherical tokamak. Figure courtesy of UKAEA [2].

pable of producing the same fusion power output, so are potentially cheaper and easier to build, making them a favourable avenue for commercial power generation. Due to the reduced plasma volume, in order to still have enough fusion reactions happening per second to meet the Lawson criterion, a higher plasma density is needed. It is plasma density that is targeted, rather than temperature as increasing temperature would move the plasma away from the temperature that maximises the cross-section of the fusion reaction. Increasing confinement time would also increase the triple product, energy confinement time has an upper limit dependent on magnetic field strength [7]. In addition, the compact design results in the magnetic field coils being closer to the plasma meaning that they require less energy to produce the same field at the plasma location. This makes it easier to achieve a higher plasma beta (ratio of plasma pressure to magnetic pressure), a common metric for the efficiency of the device [8–10].

In order to maintain the tighter configuration of these devices, it is often necessary to combine the inboard side of the toroidal field coils (the D-shaped coils in Fig. 1.1) into a single central conducting column. This leaves little room for a central solenoid, making the use of alternative current drive systems all the more important [8–10].

1.2 Turbulence in Fusion Plasmas

One of the enduring problems in achieving power generation from fusion in MCF devices is the degradation of the confinement by turbulent effects increasing cross-field transport [11,12]. In addition, turbulence in the plasma has the potential to interfere with diagnostics and power injection systems. This has made the study of plasma turbulence crucial to the success of fusion power.

Steep temperature and density gradients are one of the primary sources of this turbu-

lence, driving a number of micro-instabilities. For example, temperature and magnetic field gradients drive the electron temperature gradient (ETG) [13] and ion temperature gradient (ITG) [14] modes as well as macro-scale MHD modes. The magnetic field gradient is perpendicular to the toroidal field (which dominates the magnetic field direction) resulting in a grad-B ($\nabla\mathbf{B} \times \mathbf{B}$) drift along the flux surfaces. As hotter particles closer to the core of the device have a higher average velocity and lower collisionality, they drift faster than cooler ones closer to the edge. This means that the radial temperature gradient causes local temperature and density perturbations, generating local electric fields. These electric field perturbations give rise to local $\mathbf{E} \times \mathbf{B}$ drifts which amplify the initial perturbations [13, 14]. Due to the significantly smaller mass of electrons, the growth rate for the ETG is much larger than the ITG, though both lead to significant thermal transport as the Larmor radius scales with particle mass, meaning that the ion gyro-radius is significantly larger than the electron gyro-radius.

Another example is drift waves driven by the radial density gradient at an angle to the magnetic field. If there is a small seed perturbation in the ion density perpendicular to the density gradient, this creates a region of high and low density near each other, establishing a local electrostatic potential. The electrons respond to this faster than ions due to their lower mass, streaming along the field lines, thus generating a local electric field. The resulting $\mathbf{E} \times \mathbf{B}$ drift is then 90° out of phase with the initial density perturbation, resulting in a wave propagating perpendicular to the density gradient and magnetic field with no net transport, as the velocity and density perturbation are out of phase. However, if resistivity causes a delay in the electron response to the electrostatic potential, there is a phase shift resulting in the wave becoming unstable, generating rotating turbulent eddies that cause particle and energy transport [15, 16].

Larger scale magneto-hydrodynamic (MHD) instabilities also degrade the confinement of the device, such as the neoclassical tearing mode (NTM). These form on surfaces where the safety factor $q = m/n$ is a ratio number. Here m and n are the poloidal and toroidal mode numbers respectively. The safety factor is defined as the number of toroidal turns a field line makes for each poloidal turn. NTMs require a seed ‘island’ in the magnetic field. As particles stream along field lines, these islands flatten the pressure gradient across them, leading to a hole in the bootstrap current — a spontaneously arising current driven by collisions between particles trapped by the magnetic mirror force on the weak-field (outboard)

side of the tokamak and particles which can complete full toroidal orbits as a result of the radial pressure gradient. This perturbation to the bootstrap current, in the presence of a magnetic shear and pressure gradient, leads to a destabilizing current which reinforces and grows the seed magnetic island. NTMs induce eddy currents, degrade plasma and energy confinement, cause disruption, and slow plasma rotation [17, 18].

All of these effects, among others [11, 12], make the achievement of break-even in a fusion reactor significantly more challenging. The ability to investigate these instabilities and the turbulence they produce with diagnostics of sufficient spatial and temporal resolution is therefore of great importance, as is the ability to control or mitigate the instabilities. It is worth noting that as well as having a negative impact on the stability and confinement of MCF reactors, the turbulent fluctuations also scatter incident EM waves, particularly those with wavelength comparable to the length scale of the turbulence. This has further negative implications for the efficiency of the reactors if heating and current drive beams don't deposit their power where intended.

1.3 Microwaves in Fusion Plasmas

1.3.1 Heating and current drive

Microwaves have a frequencies in the same range as the electron cyclotron resonances (and their harmonics) in fusion plasmas, making microwave systems an effective tool to inject power into the plasma for the purposes of both global and local heating and current drive [19–22] as well as for the purposes of non-inductive start-up [23, 24]. Global heating and current drive (where numerous deposition regions are targeted to generate the desired temperature and current profile) are essential in order to reach the temperatures required for fusion to occur and to generate the poloidal field required to confine the plasma. Local heating and current drive can also be useful in stabilising instabilities such as the NTMs mentioned in Sec. 1.2.

Microwave methods for power injection have numerous advantages over other common approaches such as neutral beam injection (NBI). The antennas require little space on the vessel walls, leaving more room for breeder blankets. The gyrotrons used to generate the microwaves can be housed long distances from the reactor itself, outside of the bioshielding, meaning that maintenance can more easily be carried out without having to shut down the

reactor, and waveguides can be used to carry the microwaves to the reactor with very little loss of power. Finally, as the injected microwaves are EM waves, they can be launched in vacuum at the edge of the device before coupling efficiently with the plasma.

Two of the primary methods of power injection using microwaves are electron cyclotron resonance heating (ECRH) and electron cyclotron current drive (ECCD). These methods use EM waves in the frequency range of the electron cyclotron resonance $\omega_{ce} = eB/m_e$ or its low-order harmonics. This is the frequency at which electrons gyrate around the magnetic field lines in the plasma. This allows highly targeted and localised power deposition, due to the dependence of ω_{ce} on magnetic field strength, B , and the way B varies across a poloidal cross-section of the device [22].

However, in some over-dense plasmas where the cut-off density for a particular frequency of wave is located closer to the plasma edge than the location of the cyclotron harmonic (frequently occurring in spherical tokamaks), waves of the frequency used for ECCD and ECRH cannot propagate to the desired deposition region due to the existence of a cut-off density near the plasma edge. The exact nature of this cut-off density depends upon the mode of wave used. More details of these can be found in Sec. 2.1. While higher harmonics of the cyclotron resonance could be used, these tend to be less efficiently damped in the desired deposition region.

In these cases, electron Bernstein waves (EBW) can be used instead as they can propagate past these cut-off densities and into the targeted deposition region. These are shorter wavelength waves which only exist in hot magnetised plasmas. As they result from the coherent motion of electrons, they cannot be launched in or propagate through vacuum [25]. Instead, an EM wave must be launched which will mode convert to an EBW inside the plasma — a process which is discussed in more detail in Sec. 2.1.2.1 and Sec. 2.2. The EBW is then absorbed at a harmonic of the cyclotron resonance so can be used to heat the plasma or drive current in a similar way to ECRH and ECCD [25, 26].

In order for these methods of power injection to be effective and efficient, we must be able to accurately predict the path of microwave beams through the plasma and the area over which their power will be deposited. This is made more difficult by the turbulence present in the plasma edge, where density fluctuates on length scales comparable to microwave wavelength and the fluctuation level can reach 100% of the background density, scattering incident microwaves. Simulations are therefore an incredibly important tool in understanding

and optimising microwave heating and current drive systems [12, 27–29].

1.3.2 Diagnostics

Microwave diagnostics have good spatial and temporal resolution and can be used to both actively and passively probe fusion plasmas. Compared to diagnostics in the visible spectrum, their focusing optics are robust to neutron bombardment due to their longer wavelength not being affected by small-scale roughness.

One of the most common and important passive diagnostics is the measurement of electron cyclotron emission (ECE). Just like ECRH and ECCD, this utilises the electron cyclotron resonance. However, rather than using this for wave absorption, ECE detects the emission of electrons gyrating around the magnetic field lines in the plasma. The radiation they emit is at the electron cyclotron frequency ω_{ce} or a harmonic of it. The thermal motion of the electrons broadens this peak in frequency space so that it is not just a single frequency that is detected. When the plasma is optically thick enough (when optical depth $\tau \gg 1$), the intensity of the observed ECE approaches that of a black body, meaning it can be used to determine the electron temperature at the spatial location corresponding to a plasma density that is close to a resonance. Similarly to how the dependence of ω_{ce} on the magnetic field strength allows for targeted power deposition in the cases of ECRH and ECCD, for ECE diagnostics, this allows for targeted probing of select regions within the plasma, giving the diagnostic a spatial resolution of order 1 cm [30]. Typical ECE radiometers detect in a number of frequency bands, each corresponding to a particular location in the plasma. This allows for a time-resolved 1D line profile of the electron temperature to be determined. In order to construct a 2D temperature profile, a 1D array of detectors can be used [31].

As with the power injection methods, when the plasma is overdense such that the plasma frequency is greater than the cyclotron frequency, ECE can no longer be used as a passive probing technique as the emitted waves cannot escape the plasma to be detected. In this case, EBW emission (EBE) can be used instead. This is slightly more complex, as the EBE cannot propagate through the vacuum at the plasma edge so must mode convert to an EM wave in order to be detected. The efficiency of this process will affect the intensity of the detected EBE spectrum, so accurate knowledge of this efficiency is required. To this end, ray tracing calculation of the EBW path through the plasma to the mode conversion region are needed. These can then be used to deduce the location of the emission and calculate predicted mode

conversion efficiency in order to accurately deduce the electron temperature [25, 31].

Reflectometry is an active probing technique where an EM wave is launched and propagates into the plasma until it reaches its cut-off density, at which it is reflected back. The time between emitting and receiving a signal can be measured and used to calculate the path length of the wave and thus the location of the cut-off density. Repeating this measurement many times at different frequencies then allows the construction of the electron density profile and any fluctuations present [30, 31]. Doppler reflectometry measures the Doppler shift of the reflected signal, caused by the rotation of the plasma. This provides a measurement of the speed of the cut-off surface and can be used to find the rotation velocity of the plasma [30, 31].

With all microwave diagnostics, it is important to be able to accurately predict the beam path so that we know exactly what region of the plasma we are probing. Furthermore, with more complex diagnostics setups, comparison of diagnostic measurements to the output of simulated versions of the diagnostics can give much greater insight into the plasma properties.

1.4 Thesis Outline

The goal of this project is to better understand the propagation of microwaves through fusion plasmas, particularly in instances where plasma density is varying on length scales comparable to microwave wavelength. The basic physics of microwaves in plasmas is covered in Chapter 2. The numerical tools that can be used to investigate their propagation in fusion plasmas are covered in Chapter 3. The full-wave cold plasma code that I have developed and used for this project is described in Chapter 4. A thorough benchmark of the code is then carried out in Chapter 5 along with a study of the OX mode conversion process with varying density scale length. In Chapter 6 the full-wave cold plasma code EMIT-2D is used to simulate the broadening of ECRH beams by plasma turbulence on the DIII-D tokamak in order to compare the results of simulations to experimental measurements as an important benchmark. The code is then used to carry out a wide-ranging parametric scan in order to determine the dependence of microwave beam broadening by plasma turbulence on plasma and beam parameters which is covered in Chapter 7. Finally, the main conclusions of all of this work as well as suggestions for further work are covered in Chapter 8.

Chapter 2

Waves in Plasmas

For analysis of waves in plasmas, it is useful to define two fundamental frequencies. First, the plasma frequency,

$$\omega_{p,s} = \sqrt{\frac{Z_s n e^2}{\epsilon_0 m_s}}, \quad (2.1)$$

is the frequency at which a species, s , in the plasma responds to a small displacement. Charge separation generates an electric field resulting in a restoring force proportional to the displacement, leading to an electrostatic simple harmonic oscillator. Here Z_s is the species atomic number, n is the number density, e the charge of an electron, ϵ_0 the permittivity of free space, and m_s is the species mass.

Second, the cyclotron frequency,

$$\omega_{c,s} = \frac{q_s B}{m_s}, \quad (2.2)$$

is the frequency at which particles orbit magnetic field lines in the plasma. B is the magnetic field strength and q_s is the species charge.

These frequencies provide a normalized plasma density,

$$X = \left(\frac{\omega_{p,e}}{\omega} \right)^2, \quad (2.3)$$

and magnetic field strength,

$$Y = \frac{\omega_{c,e}}{\omega}, \quad (2.4)$$

where ω is the vacuum frequency of the wave travelling through the plasma. Here, we have

used the electrons rather than the ions because, as will be discussed later in this Chapter, it is the electrons that are most relevant for microwaves propagating through fusion plasmas. This is due to the larger mass of the ions meaning they take much longer to respond to any perturbing effect of the wave.

The derivations in this chapter draw on information covered in [16] and [32] throughout.

2.1 Waves in Cold Plasmas

In order to simplify the physics, for investigating microwave propagation in fusion plasmas, it is common to utilise the cold plasma approximation. In doing so, we neglect the thermal motion of the particles, assuming they all travel at the same velocity. This approximation applies when the wave in question propagates significantly faster than the thermal velocity in the plasma. This is the case the majority of the time for microwaves in fusion plasmas. However, it cannot be used to model wave absorption as this depends on finite Larmor radius effects.

In order to understand how waves propagate in cold plasmas, it is useful to derive a dispersion relation to tell us how wave frequency will depend on wave vector and position vector. In order to do this, we begin with Maxwell's equations,

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad (2.5)$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \frac{1}{c^2} \frac{\partial \mathbf{E}}{\partial t}, \quad (2.6)$$

and substitute in plane wave solutions of the form $\sim \exp[i(\mathbf{k} \cdot \mathbf{r} - \omega t)]$ for the perturbed part of \mathbf{E} , \mathbf{B} , and \mathbf{J} . This gives us that

$$\mathbf{k} \times \mathbf{E} = \omega \mathbf{B} \quad (2.7)$$

$$\mathbf{k} \times \mathbf{B} = -i\mu_0 \mathbf{J} - \frac{\omega}{c^2} \mathbf{E}. \quad (2.8)$$

We can then use Ohm's law to express \mathbf{J} in terms of \mathbf{E} as

$$\mathbf{J} = \underline{\underline{\sigma}} \cdot \mathbf{E}, \quad (2.9)$$

where $\underline{\underline{\sigma}}$ is the conductivity tensor, and substitute this into Eq. (2.8) to give

$$\mathbf{k} \times \mathbf{B} = -i\mu_0 \underline{\underline{\sigma}} \cdot \mathbf{E} - \frac{\omega}{c^2} \mathbf{E}. \quad (2.10)$$

This can be rewritten as

$$\mathbf{k} \times \mathbf{B} = -\frac{\omega}{c^2} \underline{\underline{\varepsilon}} \cdot \mathbf{E} \quad (2.11)$$

where

$$\underline{\underline{\varepsilon}} = \underline{\underline{I}} + \frac{i}{\varepsilon_0 \omega} \underline{\underline{\sigma}} \quad (2.12)$$

is the dielectric permittivity tensor.

From Eq. (2.7) we can see that $\mathbf{B} = \frac{1}{\omega} \mathbf{k} \times \mathbf{E}$ which can be substituted into Eq. (2.11) to get

$$\left[\mathbf{k}\mathbf{k} - k^2 \underline{\underline{I}} + \frac{\omega^2}{c^2} \underline{\underline{\varepsilon}} \right] \cdot \mathbf{E} = \underline{\underline{M}} \cdot \mathbf{E} = 0. \quad (2.13)$$

This is a general dispersion relation, meaningful solutions for which can be found where $\det(\underline{\underline{M}}) = 0$.

However, in order to solve this equation, we need to know the form of $\underline{\underline{\sigma}}$ and therefore $\underline{\underline{\varepsilon}}$. To do this, we consider the relationship between \mathbf{J} and \mathbf{E} combined with the relationship described in Eq. (2.9).

To find how current density depends on electric field strength, we start by considering the Lorentz force on an individual particle

$$\mathbf{F} = m \frac{d\mathbf{v}}{dt} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}). \quad (2.14)$$

If we take the cold plasma approximation and assume that all particles move at the same velocity rather than having a thermal distribution of velocities, then the above equation also describes the fluid velocity of the plasma. We can then use this to describe the current density $\mathbf{J}_s = n_s q_s \mathbf{v}_s$ such that

$$\frac{m_s}{n_s q_s} \frac{d\mathbf{J}_s}{dt} = q_s \mathbf{E} + \frac{1}{n_s} \mathbf{J}_s \times \mathbf{B}. \quad (2.15)$$

We will assume that the plasma is homogenous, quasi-neutral, and at rest apart from the effect of the perturbing plane-wave fields. This means that we can state that the background

electric field and current density are zero. We assume a solution of the form

$$\mathbf{E} = \mathbf{E}_1 \exp[i(\mathbf{k} \cdot \mathbf{r} - \omega t)] \quad (2.16)$$

$$\mathbf{B} = \mathbf{B}_0 + \mathbf{B}_1 \exp[i(\mathbf{k} \cdot \mathbf{r} - \omega t)] \quad (2.17)$$

$$\mathbf{J} = \mathbf{J}_1 \exp[i(\mathbf{k} \cdot \mathbf{r} - \omega t)]. \quad (2.18)$$

If we substitute these fields into Eq. (2.15) and linearise in the wave-like perturbed quantities then we get

$$\frac{m_s}{n_{0,s}q_s^2} \frac{\partial \mathbf{J}_s}{\partial t} = \mathbf{E} + \frac{B_0}{n_{0,s}q_s} \mathbf{J}_s \times \hat{\mathbf{b}}_0 \quad (2.19)$$

where B_0 is the background magnetic field strength and $\hat{\mathbf{b}}_0$ is the unit vector pointing in the direction of the background field, and n_0 is the unperturbed number density of the particle species. Terms that are 2nd order in perturbed quantities, including the convective part of the total time derivative, have been discarded. This assumption is valid within the low power limit, however, at high powers non-linear effects are observed. The threshold at which these effects become important is the subject of some study [33, 34].

We can then use the wave-like form of the fields to compute the time derivative such that

$$-i\omega \frac{m_s}{n_{0,s}q_s^2} \mathbf{J}_s = \mathbf{E} + \frac{B_0}{n_{0,s}q_s} \mathbf{J}_s \times \hat{\mathbf{b}}_0. \quad (2.20)$$

If we arbitrarily assign the background magnetic field to point along the z -direction, we can break this down into the three vector components:

$$J_{x,s} = \frac{i}{\omega} \frac{n_{0,s}q_s^2}{m_s} E_x + \frac{i}{\omega} \frac{B_0 q_s}{m_s} J_{y,s} = \frac{i}{\omega} \left(\frac{n_{0,s}q_s^2}{m_s} E_x + \omega_{c,s} J_{y,s} \right) \quad (2.21)$$

$$J_{y,s} = \frac{i}{\omega} \frac{n_{0,s}q_s^2}{m_s} E_y - \frac{i}{\omega} \frac{B_0 q_s}{m_s} J_{x,s} = \frac{i}{\omega} \left(\frac{n_{0,s}q_s^2}{m_s} E_y - \omega_{c,s} J_{x,s} \right) \quad (2.22)$$

$$J_{z,s} = \frac{i}{\omega} \frac{n_{0,s}q_s^2}{m_s} E_z. \quad (2.23)$$

We can then substitute Eq. (2.22) into Eq. (2.21) and vice versa to get

$$J_{x,s} = \frac{i}{\omega} \frac{n_{0,s}q_s^2}{m_s} \left(\frac{E_x + i(\omega_{c,s}/\omega)E_y}{1 - (\omega_{c,s}/\omega)^2} \right) \quad (2.24)$$

$$J_{y,s} = \frac{i}{\omega} \frac{n_{0,s} q_s^2}{m_s} \left(\frac{E_y - i(\omega_{c,s}/\omega) E_x}{1 - (\omega_{c,s}/\omega)^2} \right). \quad (2.25)$$

The total current will then be the sum of the contributions from the electrons and the ions.

$$\mathbf{J} = \mathbf{J}_i + \mathbf{J}_e \quad (2.26)$$

From here we will assume that for a fusion plasma primarily made up of Deuterium and Tritium $Z = 1$ so $q_i = e$ and $q_e = -e$. From quasineutrality, we also know that $n_{0,e} = n_{0,i} = n_0$. Thus

$$J_x = \frac{i}{\omega} \left[\frac{n_0 e^2}{m_i} \left(\frac{E_x + i(\omega_{c,i}/\omega) E_y}{1 - (\omega_{c,i}/\omega)^2} \right) + \frac{n_0 e^2}{m_e} \left(\frac{E_x + i(\omega_{c,e}/\omega) E_y}{1 - (\omega_{c,e}/\omega)^2} \right) \right] \quad (2.27)$$

which can be rearranged to

$$\frac{i}{\varepsilon_0 \omega} J_x = -\frac{\omega_{p,i}^2}{\omega^2} \left(\frac{E_x + i(\omega_{c,i}/\omega) E_y}{1 - (\omega_{c,i}/\omega)^2} \right) - \frac{\omega_{p,e}^2}{\omega^2} \left(\frac{E_x + i(\omega_{c,e}/\omega) E_y}{1 - (\omega_{c,e}/\omega)^2} \right). \quad (2.28)$$

Now we can use the fact that for the waves we are interested in $\omega \gg \omega_{c,i}, \omega_{p,i}$. This is due to the large mass of the ions relative to the electrons. This then lets us drop terms in $\omega_{p,i}/\omega$ and $\omega_{c,i}/\omega$ giving us

$$\frac{i}{\varepsilon_0 \omega} J_x = -\frac{\omega_{p,e}^2}{\omega^2} \left(\frac{E_x + i(\omega_{c,e}/\omega) E_y}{1 - (\omega_{c,e}/\omega)^2} \right) = -\frac{X}{1 - Y^2} E_x - \frac{iXY}{1 - Y^2} E_y. \quad (2.29)$$

Doing the same for the other components, this then leaves us with

$$\frac{i}{\varepsilon_0 \omega} J_x = -\frac{X}{1 - Y^2} E_x - \frac{iXY}{1 - Y^2} E_y \quad (2.30)$$

$$\frac{i}{\varepsilon_0 \omega} J_y = -\frac{X}{1 - Y^2} E_y + \frac{iXY}{1 - Y^2} E_x \quad (2.31)$$

$$\frac{i}{\varepsilon_0 \omega} J_z = -X E_z \quad (2.32)$$

where X and Y have been defined above in Eqs. (2.3) and (2.4).

Combining this with Eqs. (2.9) and (2.12) then gives us

$$\underline{\underline{\varepsilon}} = \begin{pmatrix} 1 - \frac{X}{1-Y^2} & -\frac{iXY}{1-Y^2} & 0 \\ \frac{iXY}{1-Y^2} & 1 - \frac{X}{1-Y^2} & 0 \\ 0 & 0 & 1 - X \end{pmatrix}. \quad (2.33)$$

To simplify the expression, we can define right and left hand polarised EM waves

$$R = 1 - \frac{X}{1+Y} \quad (2.34)$$

$$L = 1 - \frac{X}{1-Y} \quad (2.35)$$

and the electrostatic plasma oscillation

$$P = 1 - X. \quad (2.36)$$

We can then define terms based on the sum and difference of the R and L wave

$$S = \frac{1}{2}(R + L) \quad (2.37)$$

$$D = \frac{1}{2}(R - L). \quad (2.38)$$

Substituting in these definitions to Eq. (2.33) gives us

$$\underline{\underline{\varepsilon}} = \begin{pmatrix} S & -iD & 0 \\ iD & S & 0 \\ 0 & 0 & P \end{pmatrix}. \quad (2.39)$$

This can then be substituted into Eq. (2.13) to give

$$\underline{\underline{M}} = \left[\mathbf{k}\mathbf{k} - k^2 \underline{\underline{I}} + \frac{\omega^2}{c^2} \underline{\underline{\varepsilon}} \right] = \begin{pmatrix} k_x^2 - k^2 + \frac{\omega^2}{c^2} S & k_x k_y - i \frac{\omega^2}{c^2} D & k_x k_z \\ k_y k_x + i \frac{\omega^2}{c^2} D & k_y^2 - k^2 + \frac{\omega^2}{c^2} S & k_y k_z \\ k_z k_x & k_z k_y & k_z^2 - k^2 + \frac{\omega^2}{c^2} P \end{pmatrix}. \quad (2.40)$$

Using the fact that we have assigned $\mathbf{B}_0 = B_0 \hat{\mathbf{z}}$, we can set $\mathbf{k} = (k_x, k_y, k_x) = (k_\perp, 0, k_\parallel)$ without loss of generality. We can also multiply $\underline{\underline{M}}$ by $\frac{c^2}{\omega^2}$ without altering solutions to the

dispersion equation and then use the refractive index

$$\mathbf{N} = \frac{c\mathbf{k}}{\omega} = (N_x, N_y, N_z) = (N \sin \theta, 0, N \cos \theta) \quad (2.41)$$

where θ is the angle between the \mathbf{k} and \mathbf{B}_0 .

This then gives us

$$\underline{\underline{M}} \cdot \mathbf{E} = \begin{pmatrix} S - N^2 \cos^2 \theta & -iD & N^2 \sin \theta \cos \theta \\ iD & S - N^2 & 0 \\ N^2 \sin \theta \cos \theta & 0 & P - N^2 \sin^2 \theta \end{pmatrix} \cdot \mathbf{E} = 0. \quad (2.42)$$

This has meaningful solutions where $\det(\underline{\underline{M}}) = 0$ which yields

$$\tan^2 \theta = -\frac{P(N^2 - R)(N^2 - L)}{(SN^2 - RL)(N^2 - P)}. \quad (2.43)$$

There are numerous potential solutions to this equation which represent distinct modes of wave. Each of these modes has its own distinct properties such as phase velocity, resonances (where $N \rightarrow \infty$ and the wave would be absorbed), and cut-offs (where $N \rightarrow 0$ and the wave enters a region of evanescence it can't propagate through). It is worth noting that the formalism we have gone through for the cold plasma model does not accurately capture the wave behaviour at resonances, but it can be used to predict some of their locations. This is covered in more detail in Sec. 2.2.

In the following sections, we will examine the specific cases of propagation perpendicular or parallel to the magnetic field.

2.1.1 Propagation Parallel to Magnetic Field

For propagation parallel to the magnetic field, $\theta = 0$ so our dispersion relation becomes

$$\frac{P(N_{\parallel}^2 - R)(N_{\parallel}^2 - L)}{(SN_{\parallel}^2 - RL)(N_{\parallel}^2 - P)} = 0. \quad (2.44)$$

This then yields three solutions

$$N_{\parallel}^2 = R = 1 - \frac{X}{1 + Y} \quad (2.45)$$

$$N_{\parallel}^2 = L = 1 - \frac{X}{1 - Y} \quad (2.46)$$

$$P = 1 - X = 0 \quad (2.47)$$

Further information can be gained by considering the eigenvectors for each mode which can be obtained by substituting the solutions back into the matrix equation in Eq. (2.42).

Eq (2.45) corresponds to the right-hand circularly polarised electromagnetic wave with eigenvector $(E_x, iE_x, 0)$, sometimes referred to as the R-wave. Eq (2.46) corresponds to the left-hand circularly polarised electromagnetic wave with eigenvector $(E_x, -iE_x, 0)$, sometimes referred to as the L-wave.

Eq. (2.47) corresponds to an electrostatic oscillation with eigenvector $(0, 0, E_z)$.

For the two electromagnetic modes, we can see that their cut-offs (where $N \rightarrow 0$ and the wave enters a region of evanescence it can't propagate through) occur when

$$X = 1 \pm Y \quad (2.48)$$

where the '+' corresponds to the R-wave and the '-' corresponds to the L-wave. Returning to expressing things directly in terms of frequencies and rearranging results in cut-off frequencies of

$$\omega_R = \frac{1}{2} \left(\sqrt{4\omega_{p,e}^2 + \omega_{c,e}^2} + \omega_{c,e} \right) \quad (2.49)$$

$$\omega_L = \frac{1}{2} \left(\sqrt{4\omega_{p,e}^2 + \omega_{c,e}^2} - \omega_{c,e} \right) \quad (2.50)$$

where we take the positive root in each case.

When considering resonances of the waves, it is important to remember that we defined $\omega_{c,e}$ such that it is always negative for electrons, and therefore $Y = \omega_{c,e}/\omega$ is also negative. Recalling that a resonance occurs when $N \rightarrow \infty$ we can therefore see that the R-wave has a resonance at

$$Y = -1 \rightarrow \omega = -\omega_{c,e} \quad (2.51)$$

which is the electron cyclotron resonance. This means that the electric field of the wave is oscillating at the same frequency as the frequency at which electrons gyrate around magnetic field lines. Crucially, the electric field is also oscillating in the same direction as the electrons. This leads to strong absorption, though the mechanics of this are not captured in the cold plasma model.

No such resonance can be seen in the L-wave as the electric field oscillates in the opposite direction to the electrons. If we had not neglected the contributions of ions due to their large masses and slow response times, we would see a resonance at $\omega = \omega_{c,i}$, the ion cyclotron frequency. However, this is a much lower frequency and is outside the range we are interested in.

2.1.2 Propagation Perpendicular to Magnetic Field

For propagation perpendicular to the magnetic field, $\theta = \pi/2$ so $\tan^2 \theta \rightarrow \infty$ such that our cold plasma dispersion relation becomes

$$\frac{P(N_{\perp}^2 - R)(N_{\perp}^2 - L)}{(SN_{\perp}^2 - RL)(N_{\perp}^2 - P)} \rightarrow \infty. \quad (2.52)$$

This has two possible solutions

$$N_{\perp}^2 = P = 1 - X \quad (2.53)$$

$$N_{\perp}^2 = \frac{RL}{S} = \frac{(1 - X - Y)(1 - X + Y)}{1 - X - Y^2}. \quad (2.54)$$

The solution in Eq. (2.53) corresponds to the ordinary mode (O-mode), so named because of its simpler dispersion relation, similar to a wave propagating in vacuum adjusted by a term to account for the plasma density. It has eigenvector $(0, 0, E_z)$. The O-mode is a transverse, electromagnetic wave which is linearly polarised with its electric field pointing parallel to \mathbf{B}_0 .

The solution in Eq. (2.54) corresponds to the extraordinary mode (X-mode), so named because of the more complicated dispersion relation compared to the O-mode in that the propagation depends on magnetic field strength as well as plasma density. It has eigenvector $(E_x, -i\frac{S}{D}E_x, 0)$. The X-mode is an electromagnetic wave with longitudinal and transverse components resulting in an elliptical polarisation of the electric field perpendicular to \mathbf{B}_0 .

Just as for parallel propagation, these two distinct modes have distinct properties such as phase and group velocities, resonances, and cut-offs.

From the O-mode dispersion relation in equation (2.53), it can be seen that there is a cut-off where $N_{\perp} \rightarrow 0$ at

$$X = 1 \rightarrow \omega = \omega_{p,e}. \quad (2.55)$$

Hence, O-mode waves cannot propagate through a plasma where the density is such that the

plasma frequency equals the wave frequency. Throughout the rest of the thesis, the density at which this occurs is referred to as n_{crit} .

No resonance conditions for the O-mode exist in the cold plasma approximation. This is not the case in the real world, which will be discussed further in Sec. 2.2.

From the X-mode dispersion relation in equation (2.54), it can be seen that there are two cut-offs corresponding to $R = 0$ and $L = 0$. These are the same cut-off conditions that we found for R-mode and L-mode propagation parallel to the magnetic field, meaning that the X-mode has cut offs at

$$\omega = \omega_{\text{R}} \quad (2.56)$$

and

$$\omega = \omega_{\text{L}} \quad (2.57)$$

as defined in Eqs. (2.49) and (2.50). These cut-offs are independent of angle of propagation, hence they arise for waves that propagate both parallel and perpendicular to the field.

From equation (2.54) we also find a resonance (where $N_{\perp} \rightarrow \infty$) at $S = 0$ or at $1 - X - Y^2 = 0$. This corresponds to a resonant frequency at

$$\omega_{\text{UH}} = \sqrt{\omega_{\text{p,e}}^2 + \omega_{\text{c,e}}^2}. \quad (2.58)$$

This resonance contains both plasma and cyclotron frequencies and is hence known as the upper hybrid frequency as it is greater than both the plasma frequency and electron frequency. Though we have neglected the contributions of ions due to their larger masses, the inclusion of ions would yield a second resonance at the lower hybrid frequency which lies between the ion and electron cyclotron frequencies. However, this is below the frequency range of interest for us.

The X-mode dispersion relation has two distinct branches corresponding to the slow X-mode (SX) and the fast X-mode (FX), so labelled due to the magnitude of their phase velocities. The FX-mode exists in the plasma at lower densities, below the R cut-off. The SX-mode exists in the plasma at higher densities, between the upper hybrid resonance (UHR) and the L cut-off.

2.1.2.1 OX Mode Conversion

For the more generalised case of propagation at an angle θ to the background magnetic field, we can rewrite Eq. (2.43) in terms of perpendicular and parallel components of the refractive index, such that

$$N^2 = N_{\perp}^2 + N_{\parallel}^2 = 1 - \frac{2X(1-X)}{2(1-X) - Y^2 \sin^2 \theta \pm \Gamma} \quad (2.59)$$

where

$$\Gamma = (Y^4 \sin^4 \theta + 4(1-X)^2 Y^2 \cos^2 \theta)^{1/2} \quad (2.60)$$

and the ‘+’ and ‘-’ signs correspond to the O-mode and X-mode respectively. This is known as the Appleton-Hartree equation [16, 35, 36]. From this, we can see that when $X = 1$ (or writing things explicitly in terms of frequencies, when $\omega = \omega_{p,e}$) and $\theta = 0$, we get that $\Gamma = 0$, making the O-mode and X-mode degenerate. This allows for conversion from one mode to the other.

If we consider the conversion from O-mode to X-mode, the refractive index corresponding to the maximum conversion can then be calculated by substituting these values back into Eq. (2.59). If we first assume that $\theta = 0$ then we find that

$$N_{\parallel}^2 = \frac{1 - X \pm Y}{1 \pm Y}. \quad (2.61)$$

Assuming that our refractive index must be positive allows us to discard the ‘-’ in the \pm . Then, as the wave propagates into the plasma and approaches the cut-off density at $X = 1$ we find

$$N_{\parallel, \text{opt}}^2 = \frac{Y}{1 + Y}. \quad (2.62)$$

At this optimum, there is complete conversion between the two modes. Even when these conditions aren’t exactly met, mode conversion can still occur via tunnelling through an evanescent region, reducing the mode conversion efficiency. However, if the angular deviation from this optimum is too large, the wave can’t tunnel through the evanescent region so no mode conversion occurs. This results in an elliptical mode conversion window around the optimum angle of propagation.

This process is sometimes utilised to launch Electron Bernstein Waves in over dense plasmas via the OXB mode conversion.

This is the process by which an O-mode wave is launched in vacuum on the low-field side

of the tokamak. It propagates in until it reaches the O-mode cut-off density as described by Eq. (2.55). At this point, the O-mode and the SX-mode are degenerate and mode conversion between the two can take place. The SX-mode then propagates further into the plasma but turns around before it reaches the L cut-off as described in Eq. (2.50) before propagating back to the UHR as described in Eq. (2.58). At the UHR it smoothly and continuously converts to an EBW, turning back around and propagating into the plasma core where it will eventually be absorbed at a harmonic of the cyclotron resonance.

2.2 Limitations of the Cold Model

In the cold plasma model, it is assumed that electrons have zero thermal velocity resulting in the Larmor radius of their gyro-orbits of magnetic field lines being zero. Whilst this is a remarkably good approximation for microwaves propagating through fusion plasmas in a range of scenarios, there are some areas where it cannot be used.

For example, in the O-X-B mode conversion scheme discussed in Sec. 1.3.1 and Sec. 2.1.2.1, the cold plasma formalism cannot describe the EBWs, as these are plasma waves resulting from the collective coherent gyro-motion of electrons. Therefore, their description requires a warm plasma model.

Another limitation of the cold model is in the description of microwave absorption by the plasma. In Sec. 1.3.1, the use of microwaves for heating and current drive was discussed, utilising the electron cyclotron resonance. However, when considering the dispersion relations for the O-mode and X-mode given by Eqs. (2.53) and (2.54), there is no resonance at the cyclotron frequency or its harmonics. The finite Larmor radius effects introduced in the warm model allow for the absorption of the O-mode and X-mode at resonances of the cyclotron harmonic.

However, these finite temperature effects are outside the scope of this thesis. This work focuses on the propagation of the O-mode and X-mode through fusion plasmas. Their absorption is not considered and EBWs are not modelled, they are merely motivating factors in carrying out this work.

Chapter 3

Numerical Methods

3.1 Summary of Approaches

There are many different approaches to simulating the propagation of microwaves through fusion plasmas. The choice of method depends on the physics that you want to capture, weighed up against computational expense.

3.1.1 Ray Tracing

Ray tracing methods use geometrical optics to find the path of a single ray through the plasma. In order to do this, one must assume that the refractive index of the plasma varies slowly in time (compared to the wave period) and in space (compared to the wavelength). This means the Wentzel–Kramers–Brillouin (WKB) approximation can be applied, resulting in an eikonal equation to the lowest order. This eikonal is then used to construct a set of ODEs which can be solved for the ray trajectory, and rather than solving the full Maxwell equations (as is done in a full-wave code), one can instead solve the simpler ray equations [32, 37]

$$\frac{d\mathbf{r}}{dt} = -\frac{\partial\mathcal{M}/\partial\mathbf{k}}{\partial\mathcal{M}/\partial\omega} \quad ; \quad \frac{d\mathbf{k}}{dt} = \frac{\partial\mathcal{M}/\partial\mathbf{r}}{\partial\mathcal{M}/\partial\omega} \quad ; \quad \frac{d\omega}{dt} = \frac{\partial\mathcal{M}/\partial t}{\partial\mathcal{M}/\partial\omega} \quad (3.1)$$

where $\mathcal{M} = \det(\underline{M})$ is the determinant of the plasma dispersion relation matrix. This matrix could come from the cold plasma dispersion relation as shown in Eq. (2.42), the effect of temperature could be included by using a warm plasma dispersion relation, or relativistic effects could be included by using either a weakly relativistic or fully relativistic dispersion relation. One advantage of this approach is the simplicity of the ray equations makes it

more feasible to implement a more complicated dispersion relation encapsulating more of the physics than it would be to do so in a full-wave code.

The ray equations (3.1) can be solved given an initial condition for \mathbf{r} , \mathbf{k} , and ω iterating over time steps to follow the trajectory of a ray through the plasma and extracting relevant information along its path.

The ray tracing approach can be extended to beam tracing which accounts for the divergence of a launched microwave beam [38].

However, neither of these methods accounts for the scattering effect of turbulent density profiles, which broaden incident beams. Due to the expected importance of this effect, there have been attempts at including scattering via statistical models of the turbulence [27–29]. However, as the ray tracing approach is only valid when the density fluctuations occur on a length scale larger than the wavelength, or the fluctuation level is small enough that the density gradient scale length is longer than the wavelength. The wavelength of microwaves used to power injection in fusion plasmas can have wavelengths of order a millimetre up to tens of centimetres, meaning that while density scale lengths are long enough to make ray tracing approaches valid in the core of the plasma, turbulence in the plasma edge can invalidate these approximations. This limits the applicability of such an approach meaning there are some cases where the more computationally expensive full-wave approach must be taken.

3.1.2 Full-Wave Modelling

Full wave codes solve the full form of Maxwell's equations. In order to do this, an additional equation describing the dielectric response of the plasma is also needed, which is found using the plasma dispersion relation. As they are not utilising the WKB approximation, full-wave codes do not rely on the refractive index being slowly varying in space and time compared to the wavelength and wave period [39], though this is usually only needed in a narrow region near the plasma edge.

A finite element approach can be used to solve Maxwell's equations where the simulation domain is split using local approximations found by expanding the global equations [40]. However, the more common approach, and thus the focus here, is a finite difference method which solves the equations on a discretised grid. Two classes of finite difference methods exist: time domain and frequency domain [39].

Finite difference frequency domain (FDFD) methods transform Maxwell's equations to the frequency domain before discretising them [41, 42]. As there are no time steps to iterate over, the simulations are run in steady state, and the method is limited to a single frequency. Finite difference time domain (FDTD) methods iterate over the time-steps until a steady state solution is reached where the total power in the simulation domain is constant and when the wave electric field (when averaged over one wave period) is constant at each point in the domain within numerical error, and can incorporate a spectrum of frequencies in one simulation domain.

There are some scenarios where full-wave methods are not sufficient. Due to the increased complexity of the equation set compared to ray tracing, incorporating temperature effects or relativistic effects is not feasible. These effects would need to be incorporated into the plasma response equation (the derivation of which is shown in Sec. 3.2.2 and would result in an equation set that can no longer be solved using the FDTD method. This means the physics encapsulated is often limited to that which is encapsulated by the cold plasma model (for more details on these limitations, see Sec. 2.2). Near the upper hybrid resonance, a warm plasma correction can be applied to the cold plasma dielectric tensor using a low-order expansion of the full hot plasma dielectric. This allows the equations to still be solved using full-wave methods. However, moving away from the upper hybrid resonance this approximation can no longer be made, hence the need for another approach.

3.1.3 Kinetic Modelling

The codes mentioned in previous sections only calculate the wave propagating through a plasma where the electron and ion populations follow a Maxwellian distribution function (in the case of ray tracing) or use a cold plasma assumption where all the particles have the same velocity with no thermal spread. They also do not calculate the effect that the wave has on the plasma itself so do not calculate things like current drive or heating. Kinetic codes allow the inclusion of non-Maxwellian distribution functions as well as simulation of how the wave is absorbed and drives current or heating.

Particle-in-cell (PIC) codes calculate the movement of macro-particles representing many actual particles. Density and current are calculated on a grid of discrete points using the position and velocity of the macro-particles. These are then used to calculate the electric and magnetic fields on the grid, which are then used to calculate the force on the macro-particles

and thus calculate their movement. This process is then iterated over. This is usually more computationally expensive than full-wave methods when being used for equivalent problems but allows the investigation of some areas where full-wave codes are no longer valid [43–45].

Fokker-Planck codes allow the electron distribution in the plasma to be non-Maxwellian, solving the drift-kinetic equation using a Fokker-Planck collision operator. This allows for the evolution of the electron distribution function to calculate current drive [46].

3.2 The FDTD Method for Full-Wave Modelling in EMIT-3D

For the purposes of this thesis, full-wave modelling encapsulates the required physics while also being computationally inexpensive enough to run the simulations over the domain sizes and parameter ranges required to be relevant to fusion scenarios. It will therefore be the focus of the rest of this section.

Here I detail the FDTD method for full-wave modelling as it is implemented in EMIT-3D, the code that forms the basis of my PhD work.

EMIT-3D is a 3D full-wave, cold plasma code which was developed at the University of York by T. N. R. Williams [3], parallelised in 3D using MPI by M. B. Thomas [47] with further developments by D. Woodward [48].

3.2.1 The Yee Algorithm

The algorithm implemented in EMIT-3D was first proposed by Yee in 1996 [49]. It applies a centred differencing scheme to Maxwell's equations

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E}, \quad (3.2)$$

$$\frac{\partial \mathbf{E}}{\partial t} = c^2 \nabla \times \mathbf{B} - \frac{1}{\epsilon_0} \mathbf{J}. \quad (3.3)$$

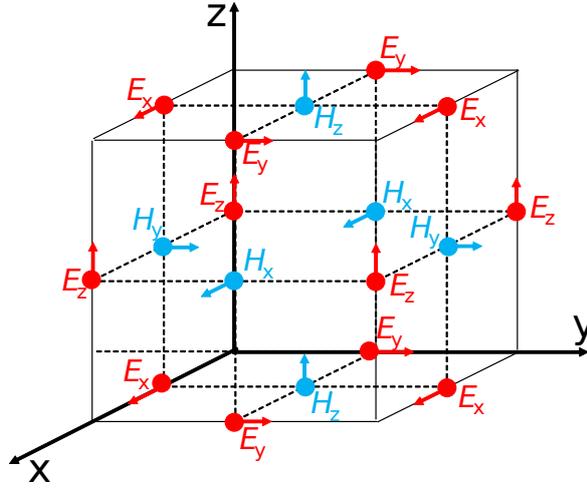


Figure 3.1: A Yee unit cell showing the grid point locations at which values are calculated. Each value is calculated using its nearest neighbours on the previous time step. The grid locations of \mathbf{E} and $\mathbf{H} = \mathbf{B}/\mu$ are shown. Figure reproduced from [3].

These can be broken down into each of their components such that

$$\frac{\partial B_x}{\partial t} = \frac{\partial E_y}{\partial z} - \frac{\partial E_z}{\partial y} \quad (3.4)$$

$$\frac{\partial B_y}{\partial t} = \frac{\partial E_z}{\partial x} - \frac{\partial E_x}{\partial z} \quad (3.5)$$

$$\frac{\partial B_z}{\partial t} = \frac{\partial E_x}{\partial y} - \frac{\partial E_y}{\partial x} \quad (3.6)$$

$$\frac{\partial E_x}{\partial t} = c^2 \left(\frac{\partial B_z}{\partial y} - \frac{\partial B_y}{\partial z} \right) - \frac{1}{\epsilon_0} J_x \quad (3.7)$$

$$\frac{\partial E_y}{\partial t} = c^2 \left(\frac{\partial B_x}{\partial z} - \frac{\partial B_z}{\partial x} \right) - \frac{1}{\epsilon_0} J_y \quad (3.8)$$

$$\frac{\partial E_z}{\partial t} = c^2 \left(\frac{\partial B_y}{\partial x} - \frac{\partial B_x}{\partial y} \right) - \frac{1}{\epsilon_0} J_z \quad (3.9)$$

As can be seen from these equations, the time evolution of \mathbf{E} depends on a spatial derivative of \mathbf{B} , and the time evolution of \mathbf{B} depends on a spatial derivative of \mathbf{E} . In order to use a centred difference approach to the FDTD, it therefore makes sense to define these two fields on displaced grids, as shown in Fig. 3.1 though in the equations, \mathbf{H} shown in the figure is replaced with $\mathbf{B} = \mu\mathbf{H}$. The current density, \mathbf{J} is defined on the same grid points as \mathbf{E} .

In this scheme, the spatial and temporal centred difference derivative of an arbitrary

parameter u can be expressed as [41]

$$\left. \frac{\partial u}{\partial x} \right|_{i,j,k}^n = \frac{u|_{i+\frac{1}{2},j,k}^n - u|_{i-\frac{1}{2},j,k}^n}{\Delta x} \quad (3.10)$$

$$\left. \frac{\partial u}{\partial t} \right|_{i,j,k}^n = \frac{u|_{i,j,k}^{n+\frac{1}{2}} - u|_{i,j,k}^{n-\frac{1}{2}}}{\Delta t} \quad (3.11)$$

where n is the time-step, i , j , and k are the grid points in x , y , and z respectively, Δx is the grid spacing, and Δt is the time-step size. This means that the actual location is given by $x = i\Delta x$ and the time is given by $t = n\Delta t$. Equivalent equations can also be written for the y and z derivatives but the grid spacing in every dimension is assumed to be uniform, and to take the same value, Δx .

The use of these centred difference derivative schemes as applied to Eqs. (3.4 – 3.9) then results in the following update equations [3, 47]:

$$B_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} = B_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n-\frac{1}{2}} - \frac{\Delta t}{\Delta x} \left[E_z|_{i,j+1,k+\frac{1}{2}}^n - E_z|_{i,j,k+\frac{1}{2}}^n - E_y|_{i,j+\frac{1}{2},k+1}^n + E_y|_{i,j+\frac{1}{2},k}^n \right] \quad (3.12)$$

$$B_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} = B_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n-\frac{1}{2}} - \frac{\Delta t}{\Delta x} \left[E_x|_{i+\frac{1}{2},j,k+1}^n - E_x|_{i+\frac{1}{2},j,k}^n - E_z|_{i+1,j,k+\frac{1}{2}}^n + E_z|_{i,j,k+\frac{1}{2}}^n \right] \quad (3.13)$$

$$B_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} = B_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n-\frac{1}{2}} - \frac{\Delta t}{\Delta x} \left[E_y|_{i+1,j+\frac{1}{2},k}^n - E_y|_{i,j+\frac{1}{2},k}^n - E_x|_{i+\frac{1}{2},j+1,k}^n + E_x|_{i+\frac{1}{2},j,k}^n \right] \quad (3.14)$$

$$\begin{aligned}
E_x|_{i+\frac{1}{2},j,k}^{n+1} &= E_x|_{i+\frac{1}{2},j,k}^n \\
&+ \frac{c^2 \Delta t}{\Delta x} \left[B_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} - B_z|_{i+\frac{1}{2},j-\frac{1}{2},k}^{n+\frac{1}{2}} - B_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} + B_y|_{i+\frac{1}{2},j,k-\frac{1}{2}}^{n+\frac{1}{2}} \right] \\
&- \frac{\Delta t}{\epsilon_0} J_x|_{i+\frac{1}{2},j,k}^{n+\frac{1}{2}}
\end{aligned} \tag{3.15}$$

$$\begin{aligned}
E_y|_{i,j+\frac{1}{2},k}^{n+1} &= E_y|_{i,j+\frac{1}{2},k}^n \\
&+ \frac{c^2 \Delta t}{\Delta x} \left[B_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} - B_x|_{i,j+\frac{1}{2},k-\frac{1}{2}}^{n+\frac{1}{2}} - B_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} + B_z|_{i-\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} \right] \\
&- \frac{\Delta t}{\epsilon_0} J_y|_{i,j+\frac{1}{2},k}^{n+\frac{1}{2}}
\end{aligned} \tag{3.16}$$

$$\begin{aligned}
E_z|_{i,j,k+\frac{1}{2}}^{n+1} &= E_z|_{i,j,k+\frac{1}{2}}^n \\
&+ \frac{c^2 \Delta t}{\Delta x} \left[B_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} - B_x|_{i-\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} - B_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} + B_x|_{i,j-\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} \right] \\
&- \frac{\Delta t}{\epsilon_0} J_z|_{i,j,k+\frac{1}{2}}^{n+\frac{1}{2}}.
\end{aligned} \tag{3.17}$$

This algorithm has a stability condition given by the Courant-Friedrichs-Lewy (CLF) number, $S = c\Delta t/\Delta x$, which must be less than some critical value, essentially requiring that a wave does not transport information more than one grid point in one timestep. The critical value for this number depends on the number of dimensions. In 1D, we require $S < 1$, in 2D it is $S < 1/\sqrt{2}$, and in 3D it is $S < 1/\sqrt{3}$ [41].

In order to time-step these equations forward, an update equation for current density is also needed.

3.2.2 The Plasma Response

In order to derive an expression for how current density evolves in time, we consider how the plasma responds to the presence of a wave [50]. To do this, we consider the electrons in the plasma as a collisionless fluid as fusion plasmas are generally low collisionality in the regimes we are considering. In addition, we neglect the contribution of ions due to their large masses resulting in slow response times compared to the wave period. Finally, as this is a cold plasma model we neglect temperature effects. The linearised fluid equation of motion of the electrons is then given by [51]

$$m_e n_e \frac{\partial \mathbf{v}_e}{\partial t} = -en_e (\mathbf{E} + \mathbf{v}_e \times \mathbf{B}_0) \tag{3.18}$$

where, as in previous sections, the background electric field and electron fluid velocity are assumed to be zero so that the only non-zero part is that perturbed by the wave. Any terms non-linear in perturbed quantities have been discarded. This means that simulations cannot directly capture the Doppler shift of waves caused by the rotation of the plasma, instead requiring such simulations to be carried out using a series of snapshots of the plasma to calculate the expected Doppler shift.

Using the previous assumption that $\omega \gg \omega_{p,i}$ meaning that the contribution to current from the ions is negligible we can write that $\mathbf{J} = -en_e \mathbf{v}_e$ along with the previously given definitions for plasma and cyclotron frequency, this can then be rewritten as

$$\frac{\partial \mathbf{J}}{\partial t} = \epsilon_0 \omega_{p,e}^2 \mathbf{E} - \omega_{c,e} \mathbf{J} \times \hat{\mathbf{b}}_0. \quad (3.19)$$

When needed, an artificial collisional damping term of $-\nu \mathbf{J}$ can be included on the left-hand side of this equation, such as when dealing with simulation of the OX-mode conversion where a lack of power dissipation mechanism can result in nonphysical power build-up at the upper hybrid resonance which would, in a warm plasma code, be carried away by electron Bernstein waves. This collision frequency is artificial and non-physical, with the value chosen to be sufficient to damp any build up of power at the UHR while not being so high as to damp the rest of the wave significantly.

The discretisation of this for the update equations in the code is best understood by rewriting it as a matrix equation.

$$\frac{\partial \mathbf{J}}{\partial t} = P \mathbf{J} + \epsilon_0 \omega_{p,e}^2 \mathbf{E} \quad (3.20)$$

where

$$P = \begin{pmatrix} 0 & -\hat{b}_z \omega_{c,e} & \hat{b}_y \omega_{c,e} \\ \hat{b}_z \omega_{c,e} & 0 & -\hat{b}_x \omega_{c,e} \\ -\hat{b}_y \omega_{c,e} & \hat{b}_x \omega_{c,e} & 0 \end{pmatrix}. \quad (3.21)$$

The resulting update equations are derived in [3] and are

$$J_x|_{i+\frac{1}{2},j,k}^{n+\frac{1}{2}} = \Theta_{1,1}J_x|_{i+\frac{1}{2},j,k}^{n-\frac{1}{2}} + \Theta_{1,2}J_y|_{i,j+\frac{1}{2},k}^{n-\frac{1}{2}} + \Theta_{1,3}J_z|_{i,j,k+\frac{1}{2}}^{n-\frac{1}{2}} \quad (3.22)$$

$$+ \omega_{p,e}^2 \left(\Xi_{1,1}E_x|_{i+\frac{1}{2},j,k}^n + \Xi_{1,2}E_y|_{i,j+\frac{1}{2},k}^n + \Xi_{1,3}E_z|_{i,j,k+\frac{1}{2}}^n \right)$$

$$J_y|_{i,j+\frac{1}{2},k}^{n+\frac{1}{2}} = \Theta_{2,1}J_x|_{i+\frac{1}{2},j,k}^{n-\frac{1}{2}} + \Theta_{2,2}J_y|_{i,j+\frac{1}{2},k}^{n-\frac{1}{2}} + \Theta_{2,3}J_z|_{i,j,k+\frac{1}{2}}^{n-\frac{1}{2}} \quad (3.23)$$

$$+ \omega_{p,e}^2 \left(\Xi_{2,1}E_x|_{i+\frac{1}{2},j,k}^n + \Xi_{2,2}E_y|_{i,j+\frac{1}{2},k}^n + \Xi_{2,3}E_z|_{i,j,k+\frac{1}{2}}^n \right)$$

$$J_z|_{i,j,k+\frac{1}{2}}^{n+\frac{1}{2}} = \Theta_{3,1}J_x|_{i+\frac{1}{2},j,k}^{n-\frac{1}{2}} + \Theta_{3,2}J_y|_{i,j+\frac{1}{2},k}^{n-\frac{1}{2}} + \Theta_{3,3}J_z|_{i,j,k+\frac{1}{2}}^{n-\frac{1}{2}} \quad (3.24)$$

$$+ \omega_{p,e}^2 \left(\Xi_{3,1}E_x|_{i+\frac{1}{2},j,k}^n + \Xi_{3,2}E_y|_{i,j+\frac{1}{2},k}^n + \Xi_{3,3}E_z|_{i,j,k+\frac{1}{2}}^n \right)$$

where Θ and Ξ have the following meanings

$$\Theta_{i,j} = \hat{b}_i\hat{b}_j[1 - \cos(\Delta t\omega_{c,e})] - \epsilon_{i,j,k}\hat{b}_k \sin(\Delta t\omega_{c,e}) + \delta_{i,j} \cos(\Delta t\omega_{c,e}) \quad (3.25)$$

$$\Xi_{i,j} = \hat{b}_i\hat{b}_j\alpha - \epsilon_{i,j,k}\hat{b}_k\beta + \delta_{i,j}\gamma \quad (3.26)$$

$$\alpha = \omega_{c,e}^2 - \omega_{c,e} \sin(\Delta t\omega_{c,e}) \quad (3.27)$$

$$\beta = \omega_{c,e} - \omega_{c,e} \cos(\Delta t\omega_{c,e}) \quad (3.28)$$

$$\gamma = \omega_{c,e} \sin(\Delta t\omega_{c,e}) \quad (3.29)$$

and where $\epsilon_{i,j,k}$ is the Levi-Cevita symbol and $\delta_{i,j}$ is the Kronecker delta symbol.

Eqs. (3.22 - 3.24) along with Eqs. (3.12 - 3.17) were the update equations implemented in the code EMIT-3D as I inherited it. This scheme altered the stability condition for the FDTD scheme slightly and a full Von Neuman analysis was carried out in [3] which found that this method would be stable as long as the Courant-Friedrichs-Lewy (CFL) condition was obeyed.

For clarity, it is worth pointing out here that though the updated current density is used to calculate the other fields, the plasma density profile used in the code itself does not evolve in time.

3.2.3 The Antenna

In order to excite the waves in the code, a soft source antenna array is used to generate a 3D Gaussian microwave beam. This means that the electric field from the antenna is added

to what is already present, rather than being set to a hard value. The electric field excited at the antenna is purely in the x -direction with the wave propagating in the z -direction, allowing the excitation of pure O-mode or pure X-mode (if launched in vacuum) by setting the background magnetic field to point in the x - or y -direction respectively. In the context of a tokamak, this is approximately equivalent to z being the radial direction and x or y being the toroidal direction depending on how the background magnetic field is set.

The expression for the electric field excited by the antenna is then given by the standard equation for a 3D Gaussian beam,

$$E_x(r, z) = E_0 \frac{w_0}{w(z)} \exp \left[\frac{-r^2}{w(z)^2} - i \left(kz + k \frac{r^2}{2R(z)} - \psi(z) \right) \right] \sin(\omega t), \quad (3.30)$$

where the following definitions are used

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R} \right)^2} \quad (3.31)$$

$$R(z) = z \left(1 + \left(\frac{z_R}{z} \right)^2 \right) \quad (3.32)$$

$$\psi(z) = \arctan \left(\frac{z}{z_R} \right) \quad (3.33)$$

$$z_R = \frac{\pi w_0^2}{\lambda}. \quad (3.34)$$

In these expressions, z is the distance along the beam path, k is the wave vector in direction of propagation (z) at the centre of the beam at it's waist location, r is the distance from the beam centre, and w_0 is the beam waist radius at the focal point. Usually, the focal point is set to be at the antenna, at $z = 0$, but it can be set at any point along the beam forwards or backwards. The beam waist radius along the beam line is given by $w(z)$ and is defined as the distance from the beam centre at which the electric field amplitude falls to $1/e$ of its peak value corresponding to power decreasing by a factor of $1/e^2$. The radius of curvature of the wave fronts is given by $R(z)$, and $\psi(z)$ is the Guoy phase, an additional phase term to account for the apparent increase in wavelength and phase velocity near the waist. The Rayleigh range, z_R , is the distance along the direction of propagation of a beam from its focal point to the point at which its cross-sectional area doubles (or the point at which its beam waist radius is multiplied by $\sqrt{2}$).

In order to ensure that only the desired frequency is launched (as opposed to the multiple

frequencies which would be generated by an instantaneous ‘switch-on’), E_x at the antenna is ramped up to the value given in Eq. (3.30) over a period of five wave-periods.

The code also includes the option of having the beam propagate at any angle. As the code is in Cartesian coordinates, in order to use this Gaussian beam expression in 3D we can express r and z in spherical polar coordinates in terms of their x , y , and z components in Cartesian coordinates and angle θ from the z -axis and angle ϕ from the xz -plane. This gives

$$r^2 = (x - x_0)^2(\sin \phi + \cos \theta \cos \phi)^2 + (y - y_0)^2(\cos \phi + \cos \theta \sin \phi)^2 \quad (3.35)$$

$$z = (x - x_0) \sin \theta \cos \phi + (y - y_0) \sin \theta \sin \phi + z_0. \quad (3.36)$$

However, when propagating at an angle to the magnetic field which is not purely perpendicular, it is no longer possible to excite a pure O-mode or a pure X-mode in the code. This can be corrected by allowing for an elliptical polarisation of the beam rather than a purely linear polarisation, which I implemented as described in Chapter 4. This is a required behaviour to be able to simulate as antennas are often designed to launch as close to pure O-mode or X-mode as possible, so the code needs to be able to do the same if it is to be useful in predictions and comparisons to experiment.

It is worth noting that the soft source antenna described in the previous section exists inside the main simulation domain so is not within the boundary layer (or at the boundary itself) as described in the following section. The location of the antenna within the simulation domain can be set by the user, but should always be in vacuum.

3.2.4 Boundary Conditions

In order to solve the update equations on the edge grid points of the simulation domain, boundary conditions are needed. In EMIT-3D, damping or absorbing boundaries are used. These reduce the amplitude of the wave so that it is approaching zero at the edge of the boundaries. The edge grid points are then set exactly to zero. Any reflections at the edge will have travelled through the absorbing boundaries twice before they re-enter the simulation domain, so reflected power is negligible.

This is implemented in the code by having a boundary layer of thickness $d_{\text{bound}} = 3\lambda_0$ where λ_0 is the vacuum wavelength. In this boundary, the wave electric field is then multi-

plied by the parabolic function

$$D(r) = 1 + \frac{13}{T} \left(\frac{r - d_{\text{bound}}}{d_{\text{bound}}} \right)^3 \quad (3.37)$$

where T is the wave period and r is the distance travelled into the boundary such that $r \leq d_{\text{bound}}$.

A thorough testing and benchmark of these boundaries can be found in [47] where it is shown that any reflections from either the gradient of the damping function or from the last grid point are negligible.

Chapter 4

Code Development: EMIT

The code I inherited at the beginning of my PhD is described in Sec. 3.2. Throughout the course of my PhD I continued to develop the code to suit my needs and fix bugs that I found as well as writing a 2D version. The changes I made are described here, and the version of EMIT used for my simulations had these changes implemented in it.

The current version of the code with my updates is not currently freely available, but is accessible via invitation on a private git repository owned by my supervisor (Roddy Vann).

4.1 EMIT-3D

Whilst I primarily worked with a 2D version of the code that will be described in Sec. 4.2, I also maintained and developed EMIT-3D so that it could be used for future work. This section describes the changes made to the code, which were also implemented in the 2D version that I wrote.

4.1.1 Change to Update Equations

It was found that when simulating the propagation of an X-mode beam, there was a slight asymmetry in the propagation resulting in the beam travelling at a small angle away from the perpendicular.

This was determined to be as a result of the update equations for current density, as given in Eq. (3.22 - 3.24). If we consider just the update equation for J_x restated here for

convenience

$$\begin{aligned}
J_x|_{i+\frac{1}{2},j,k}^{n+\frac{1}{2}} = & \Theta_{1,1}J_x|_{i+\frac{1}{2},j,k}^{n-\frac{1}{2}} + \Theta_{1,2}J_y|_{i,j+\frac{1}{2},k}^{n-\frac{1}{2}} + \Theta_{1,3}J_z|_{i,j,k+\frac{1}{2}}^{n-\frac{1}{2}} \\
& + \omega_{p,e}^2 \left(\Xi_{1,1}E_x|_{i+\frac{1}{2},j,k}^n + \Xi_{1,2}E_y|_{i,j+\frac{1}{2},k}^n + \Xi_{1,3}E_z|_{i,j,k+\frac{1}{2}}^n \right)
\end{aligned} \tag{4.1}$$

we can see that the updated value depends on J_y , J_z , E_y , and E_z as defined at slightly different locations due to the discretised grid, leading to a gradual drift of the beam centre over the course of a simulation. This asymmetry would tend to zero as the grid resolution became infinitesimal. However, to avoid setting a stringent resolution limit in the code and drastically increasing the computational expense, another solution was needed.

Instead, an average value of each of the components was taken over the four closest grid points such that the new update equation became

$$\begin{aligned}
J_x|_{i+\frac{1}{2},j,k}^{n+\frac{1}{2}} = & \Theta_{1,1}J_x|_{i+\frac{1}{2},j,k}^{n-\frac{1}{2}} \\
& + \Theta_{1,2}\frac{1}{4}\left(J_y|_{i+1,j+\frac{1}{2},k}^{n-\frac{1}{2}} + J_y|_{i,j+\frac{1}{2},k}^{n-\frac{1}{2}} + J_y|_{i+1,j-\frac{1}{2},k}^{n-\frac{1}{2}} + J_y|_{i,j-\frac{1}{2},k}^{n-\frac{1}{2}}\right) \\
& + \Theta_{1,3}\frac{1}{4}\left(J_z|_{i+1,j,k+\frac{1}{2}}^{n-\frac{1}{2}} + J_z|_{i,j,k+\frac{1}{2}}^{n-\frac{1}{2}} + J_z|_{i+1,j,k-\frac{1}{2}}^{n-\frac{1}{2}} + J_z|_{i,j,k-\frac{1}{2}}^{n-\frac{1}{2}}\right) \\
& + \omega_{p,e}^2 \left[\Xi_{1,1}E_x|_{i+\frac{1}{2},j,k}^n \right. \\
& + \Xi_{1,2}\frac{1}{4}\left(E_y|_{i+1,j+\frac{1}{2},k}^n + E_y|_{i,j+\frac{1}{2},k}^n + E_y|_{i+1,j-\frac{1}{2},k}^n + E_y|_{i,j-\frac{1}{2},k}^n\right) \\
& \left. + \Xi_{1,3}\frac{1}{4}\left(E_z|_{i+1,j,k+\frac{1}{2}}^n + E_z|_{i,j,k+\frac{1}{2}}^n + E_z|_{i+1,j,k-\frac{1}{2}}^n + E_z|_{i,j,k-\frac{1}{2}}^n\right) \right]
\end{aligned} \tag{4.2}$$

and similarly for J_y and J_z .

4.1.2 Elliptical polarisation

In order to simulate a pure O-mode propagating at an oblique angle to the magnetic field, an elliptical polarisation is required. This is needed for simulations of the OX-Mode conversion process as well as potential Doppler back-scattering simulations.

The antenna remains a soft source and still launches a 3D Gaussian beam, but rather than launching a wave linearly polarised only in the x -direction it can now launch a wave with elliptical polarisation in the x - and y - and z -directions.

From [52], we know that in order to launch a pure O-mode travelling at an angle φ to the background field which points in the z -direction, the ratio of the axes in the polarisation ellipse must be

$$R_{O_{xy}} = \frac{E_x}{E_y} = -\frac{Y^2 \sin^2 \varphi + \sqrt{Y^2 \sin^4 \varphi + 4 \cos^2 \varphi}}{2 \cos \varphi} \quad (4.3)$$

$$R_{O_{xz}} = \frac{E_x}{E_z} = -\frac{\cos \varphi}{\sin \varphi}. \quad (4.4)$$

Translating these expressions into EMIT-3D's coordinate system, $\varphi = \frac{\pi}{2} - \theta$ where θ is the angle of beam propagation from the z -axis in the xz -plane. This then gives us these new expressions for the wave excited by the antenna

$$\begin{aligned} E_{\text{Gauss}} &= E_0 \frac{w_0}{w(z)} \exp \left[\frac{-r^2}{w(z)^2} \right] \\ E_x &= E_{\text{Gauss}} \cos \left(kz + k \frac{r^2}{2R(z)} - \psi(z) - \omega t \right) \\ E_y &= \frac{1}{R_{O_{xy}}} E_{\text{Gauss}} \cos \left(kz + k \frac{r^2}{2R(z)} - \psi(z) - \omega t \right) \\ E_z &= -\frac{1}{R_{O_{xz}}} E_{\text{Gauss}} \sin \left(kz + k \frac{r^2}{2R(z)} - \psi(z) - \omega t \right) \end{aligned} \quad (4.5)$$

where we can see that the oscillations in z are 90° out of phase with those in x and y . Note that this assumes the beam is travelling only in the xz -plane with the background magnetic field aligned in the x -direction. The beam is still launched with a Gaussian profile in the y -direction, it is just that its propagation direction does not have a y -component.

This now allows the launch of pure O-mode propagating at an angle to the magnetic field. This will be particularly useful for any simulations of OX mode conversion.

4.2 EMIT-2D

Many of the scenarios of interest for my PhD were essentially 2D problems, as turbulence is extended along magnetic field lines in fusion plasmas, making most quantities close to uniform in the third dimension. Previous study has shown that 2D and 3D codes achieve the same results for such problems [53]. As such, it made sense to develop a 2D code. This gives me the ability to use larger simulation domains, simulating the power injection beams all the way into the region where they would be absorbed and using realistic microwave beam

waists. It also gives me the ability to carry out more simulations, such as those required for large-scale parameter scans.

To this end, I wrote a 2D version of EMIT-3D named EMIT-2D.

4.2.1 Update Equations in 2D

In order to create a 2D version of the code, it is simply assumed that everything is uniform in the third dimension. Because the original code excited a wave which was linearly polarised in the x -direction and travelling in the z -direction, it made sense to eliminate the y -direction. This means that any derivatives with respect to y become zero, resulting in these new update equations:

$$B_x \Big|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} = B_x \Big|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n-\frac{1}{2}} - \frac{\Delta t}{\Delta x} \left[-E_y \Big|_{i,j+\frac{1}{2},k+1}^n + E_y \Big|_{i,j+\frac{1}{2},k}^n \right] \quad (4.6)$$

$$B_y \Big|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} = B_y \Big|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n-\frac{1}{2}} - \frac{\Delta t}{\Delta x} \left[E_x \Big|_{i+\frac{1}{2},j,k+1}^n - E_x \Big|_{i+\frac{1}{2},j,k}^n - E_z \Big|_{i+1,j,k+\frac{1}{2}}^n + E_z \Big|_{i,j,k+\frac{1}{2}}^n \right] \quad (4.7)$$

$$B_z \Big|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} = B_z \Big|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n-\frac{1}{2}} - \frac{\Delta t}{\Delta x} \left[E_y \Big|_{i+1,j+\frac{1}{2},k}^n - E_y \Big|_{i,j+\frac{1}{2},k}^n \right] \quad (4.8)$$

$$\begin{aligned}
E_x|_{i+\frac{1}{2},j,k}^{n+1} &= E_x|_{i+\frac{1}{2},j,k}^n \\
&+ \frac{c^2 \Delta t}{\Delta x} \left[-B_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} + B_y|_{i+\frac{1}{2},j,k-\frac{1}{2}}^{n+\frac{1}{2}} \right] \\
&- \frac{\Delta t}{\epsilon_0} J_x|_{i+\frac{1}{2},j,k}^n
\end{aligned} \tag{4.9}$$

$$\begin{aligned}
E_y|_{i,j+\frac{1}{2},k}^{n+1} &= E_y|_{i,j+\frac{1}{2},k}^n \\
&+ \frac{c^2 \Delta t}{\Delta x} \left[B_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} - B_x|_{i,j+\frac{1}{2},k-\frac{1}{2}}^{n+\frac{1}{2}} - B_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} + B_z|_{i-\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} \right] \\
&- \frac{\Delta t}{\epsilon_0} J_y|_{i,j+\frac{1}{2},k}^n
\end{aligned} \tag{4.10}$$

$$\begin{aligned}
E_z|_{i,j,k+\frac{1}{2}}^{n+1} &= E_z|_{i,j,k+\frac{1}{2}}^n \\
&+ \frac{c^2 \Delta t}{\Delta x} \left[B_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} - B_x|_{i-\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} \right] \\
&- \frac{\Delta t}{\epsilon_0} J_z|_{i,j,k+\frac{1}{2}}^n.
\end{aligned} \tag{4.11}$$

As the plasma response update equations for current density, \mathbf{J} , do not contain any spatial derivatives, they remain unchanged. These are the equations as implemented in the code.

4.2.2 Antenna in 2D

As all derivatives in the 3rd dimension (in the y -direction) are assumed to be zero, this means that we now have plane-wave solutions in this dimension, with a 2D Gaussian beam in the xz -plane (where by 2D, I mean it has a Gaussian profile in the 1st dimension and propagates in the 2nd dimension). The equations for the electric field of the beam (either E_x for a linear polarisation as given by Eq. (3.30) or a combination of E_x and E_y for the elliptical polarisation as given by Eqs. (4.5)) are unchanged. What is different is the definitions of r and z appearing in these equations.

The 2D code still allows for propagation out of the xz -plane at an angle of ϕ , though this is almost always set to zero, as well as a propagation angle of θ from the z -axis. We can therefore express the beam coordinate in terms of these angles and Cartesian coordinates x and z :

$$r^2 = (x - x_0)^2 (\sin \phi + \cos \theta \cos \phi)^2 \tag{4.12}$$

$$z = (x - x_0) \sin \theta \cos \phi + z_0. \tag{4.13}$$

The option of elliptical polarisation, as described in Sec. 4.1.2, was also implemented in EMIT-2D.

4.2.3 Parallelisation in OpenMP

EMIT-3D was initially parallelised using MPI due to the large computational cost of a 3D code. As a 2D code is significantly computationally cheaper, and due to the growing numbers of cores with shared memory on a single node of HPC facilities, I decided that OpenMP parallelisation would be sufficient for EMIT-2D. This is a reasonable choice as the majority of the computational expense comes from performing ‘for loops’ over the grid points which can be very efficiently parallelised in OpenMP. It is also more efficient due to the shared memory meaning there is no need for additional ‘ghost cells’ in the simulation, as are present in EMIT-3D and described in [47] and also no need for taking time to pass messages between nodes at each timestep. In addition, it makes the input and output file system significantly easier to use, allowing for input density and magnetic field to come from a single file, and the output quantities to be output to a single file, rather than one file per node.

The speedup from this parallelisation can be seen in Fig. 4.1 where near perfect speed-up is observed until we approach 50 cores. The simulations conducted in this test utilised a grid of $5000 \times 5000 = 250 \times 10^3$ and $10000 \times 10000 = 1 \times 10^6$ as these were comparable parameters to those I anticipated would be required for many of my simulations. The size of the grid was considered to be the most important factor here, seeing as the parallelisation occurs over the spatial grid rather than the number of timesteps. Both simulations were run for 5000 timesteps, which is likely less than a real simulation would need to be run for to reach steady state. On one core, the simulations with 250×10^3 grid-points took 8.15 hours of wall time and on four cores, the simulation with 1×10^6 grid-points took 7.91 hours of wall time.

4.2.4 Convergence Testing, Choice of Timestep and Resolution, and Numerical Error

As discussed in Sec. 3.2, one constraint placed on the time-step is the CFL stability condition. This sets the timestep in relation to the spatial resolution, meaning that decreasing the spatial resolution (having more Yee cells per vacuum wavelength so a smaller spatial step

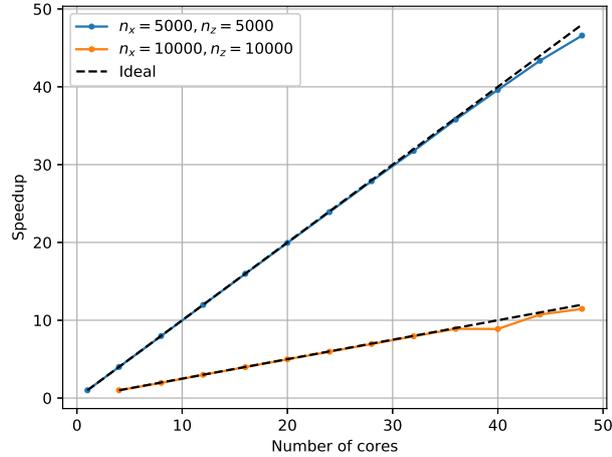


Figure 4.1: Speedup of EMIT-2D on the MARCONI supercomputer [4] for two different simulation domain sizes shown by the blue and orange data points. The black dashed line represents the ideal speedup where a doubling of cores leads to a halving of wall time. Near perfect speedup can be seen with only slight deviation after 40 cores is exceeded. Running for 5000 timesteps, the runs on the smallest number of cores (one and four cores respectively) took approximately 8 hours of wall time.

between Yee cells) also decreases the timestep. It is also possible to simply decrease the timestep while keeping the resolution the same by changing the CFL number. Both these changes improve the numerical accuracy of the code at the cost of computational time. It is therefore desirable to choose the resolution and timestep that is accurate enough whilst also as computationally cheap as possible. This can be achieved through convergence testing, decreasing either resolution or timestep until convergence (within some limit) of the end numerical result is achieved.

I did not carry out formal convergence testing with EMIT-2D, as I was using the same algorithm in the same parameter space as covered by previous PhD students [3,47] allowing me to use their numerical parameters as a starting point. In fact, the most stringent constraint on resolution was found to be the structure size in the density rather than numerical dispersion. For most cases, a resolution of 25 Yee Cells per wavelength was found to be sufficient to capture the smallest structure size, though in some cases, this was increased up to 50 Yee Cells per wavelength. This was determined through a trial and error approach, checking that a line readout of density throughout the simulation looked smooth rather than having significant discrete steps. The numbers used for each simulation are given in the relevant results chapters.

For a 2D code, the stability limit requires that $CFL < 1/\sqrt{2} \approx 0.707$. For the majority of simulations, a $CFL = 0.5$ was used in order to be accurate enough, within the stability limit, and to be convenient such that a time period could be encapsulated within an integer number of timesteps.

In order to verify that this was, indeed, accurate enough and to estimate numerical error, results for $CFL = 0.5$ were compared to an equivalent simulation with $CFL = 0.25$. This comparison was carried out for simulation set-ups which match the three scenarios in Chapter 6 and for the base case of the parameter scans as defined in Chapter 7. The example shown here is from the parameter scan base case. The primary end result of interest in this thesis is beam broadening factor, so this is the end result that was compared. This end result is reached by simulating a beam propagating through an ensemble of turbulent density profiles and calculating the ensemble average beam profile at the backplane of the simulation. A Gaussian is then fit to this ensemble average to find its full width at half maximum (FWHM). To calculate the broadening, this is compared to the FWHM of a Gaussian beam which propagated through an equivalent background plasma with no turbulence present. The results of this comparison can be seen in Fig. 4.2. In this case, a halving of the numerical time-step has resulted in a 0.1 % numerical error. This is significantly smaller than the statistical uncertainty associated with the ensemble averaging, so can be neglected. A similar result was seen for all such comparisons, with numerical error remaining as at least one order of magnitude smaller than other sources of uncertainty.

4.2.5 Benchmarking

As an initial benchmark of the 2D code, I verified that it could reproduce the O-mode and X-mode dispersion relations as described in Eqs. (2.53) and (2.54) respectively.

In order to do this, I carried out a series of simulations with a homogeneous plasma with the background magnetic field set at $Y = \omega_{c,e}/\omega = 0.5$. The direction of the magnetic field was aligned to launch pure O-mode or X-mode. The uniform plasma density was varied across the simulations. For the O-mode, it was scanned between $X = (\omega_{p,e}/\omega)^2 = 0.0 - 0.96$ so as to approach the O-mode cut-off as described in Eq. (2.55) to ensure the code reproduced the behaviour of the wavelength tending to infinity at this point. For the X-mode, the density was scanned between $X = 0.0 - 1.36$, ensuring that a resonance as described in Eq. (2.58) and two cut-offs as described in Eqs. (2.56) and (2.57) were included within the range to

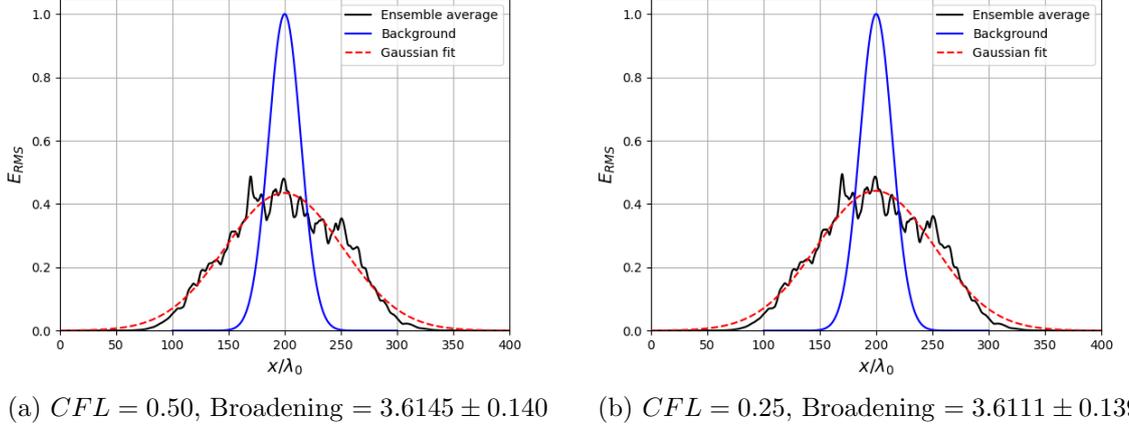


Figure 4.2: Comparison of beam broadening results when halving the time-step. The ensemble average shows the average beam after travelling through an ensemble of turbulence profiles. The background shows the beam at the same location in the simulation domain after propagating through an equivalent background plasma with no turbulence present. The broadening is found by comparing the FWHM of the ensemble average and the background beam. The uncertainty values come from the standard error on the mean of the ensemble average. From this comparison, we can see that halving the timestep decreases the broadening by 0.0034, equating to a numerical error or around 0.1 %, which is significantly smaller than the uncertainty on the result. The numerical uncertainty due to discretisation can therefore be neglected as insignificant compared to other sources of uncertainty.

verify that EMIT-2D reproduced the behaviour of the wavelength tending to infinity at the cut-offs and to zero at the resonance.

The simulation domain was $10\lambda_0 \times 10\lambda_0$, and the wavelength was calculated by finding the average peak-to-peak distance within the plasma, with the data being collected after 100 wave periods when steady-state had been reached. The resolution used was 30 Yee cells per vacuum wavelength with a CFL number of 0.5. However, near the X-mode resonance as the wavelength tended to zero in the plasma (for density values of $X = 0.8, 0.84, 0.88$, and 0.92), a greater resolution of 60 Yee cells per vacuum wavelength was used.

The results are shown in Fig. 4.3 where good agreement between the code and analytical theory is seen.

Further benchmarking by comparing to other codes can be found in Chapter 5 and a comparison to experiments can be found in Chapter 6.

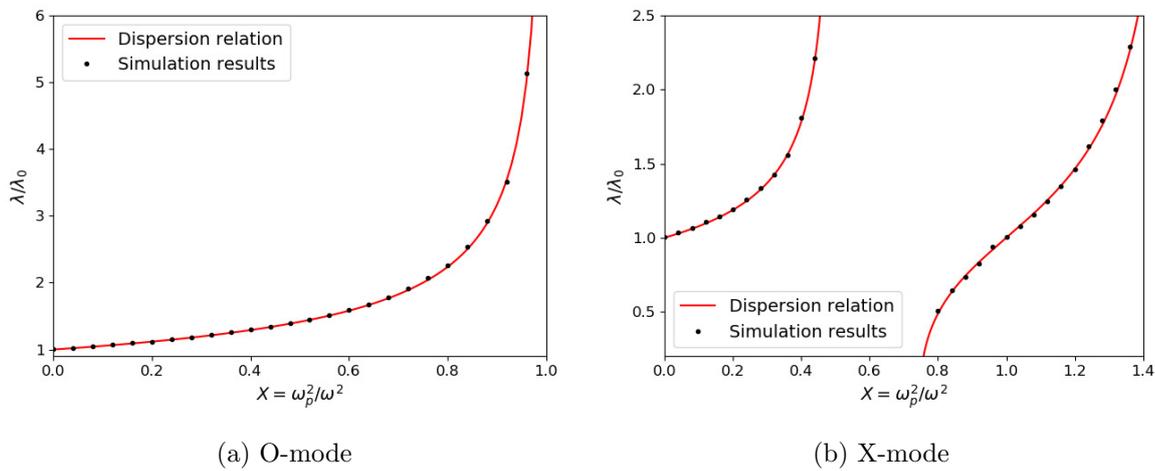


Figure 4.3: Relationship between normalised plasma density and wavelength of (a) an O-mode (left) and (b) an X-mode (right) wave propagating through the plasma. The black dots are from simulations in EMIT-2D, and the red lines are calculated directly from the dispersion relations given by Eqs. (2.53) and (2.54) respectively.

Chapter 5

Code Comparison and OX Mode Conversion Study

This work was the subject of a presentation by a collaborator at the 21st Joint Workshop on Electron Cyclotron Emission and Electron Cyclotron Resonance Heating (EC21). It was also the subject of a proceedings paper which can be found at EPJ Web of Conferences **277**, 01010 (2023) DOI: <https://doi.org/10.1051/epjconf/202327701010> [54]. As such, much of the following chapter is reproduced directly. It is also worth noting that I conducted the simulations in EMIT-2D while my collaborators independently conducted their simulations using their own codes. We then compared our results together.

This work was conducted as part of the project to install gyrotrons on MAST-U. MAST-U is a spherical tokamak [55] with a novel divertor system. As a spherical tokamak, it is often run in over-dense configurations, meaning that the cut-offs for the O-mode and X-mode waves occur further out in the plasma than the location of the cyclotron frequency where any such wave would be absorbed. This means that in order to inject power into the plasma via microwaves, the OXB mode conversion scheme as described in Sec. 2.1.2.1 must be used.

In preparation for this project, a benchmark of a suite of codes (including EMIT-2D) was performed to ensure their reliability, and an initial investigation into the effect of density scale length on the OX mode conversion efficiency was carried out.

As the codes are cold-plasma codes, only the OX mode conversion was simulated, rather than the full OXB scheme.

Scenario	Plasma	z_{ant}	x_{ant}	f_0	w_0	z_0	θ
1	vacuum (no B-field)	0	0.2 m	28 GHz	$4\lambda_0$	0.0 m	0°
2	vacuum (no B-field)	0	0.2 m	28 GHz	$4\lambda_0$	0.2 m	0°
3	vacuum (no B-field)	0	0.2 m	28 GHz	$4\lambda_0$	0.0 m	30°
4	plasma (no B-field)	0	0.2 m	28 GHz	$4\lambda_0$	0.0 m	30°
5	magnetised plasma	0	0.2 m	28 GHz	$4\lambda_0$	0.0 m	30°
6	magnetised plasma	0	0.2 m	28 GHz	$4\lambda_0$	0.0 m	θ_{opt}

Table 5.1: A summary of the scenarios used to benchmark the codes against each other. The parameters x_{ant} and z_{ant} define the location of the antenna, and θ is the angle of propagation of the beam to the z -axis. All other symbols have their previously given meanings.

5.1 Code Comparison

5.1.1 Simulation Set-Up

The codes compared were EMIT-2D (as described in Sec. 4.2), a 2.5D full-wave code called IPF-FDMC [56], and a Fourier Full-Wave code (FFW) [57]. The other codes were developed and run by my collaborators.

Though some of the minor details are different, such as how the boundaries are handled or the implementation of the algorithm, EMIT-2D and IPF-FDMC work using the same FDTD approach as previously described. The Fourier Full-Wave code also solves the same set of equations (Maxwell’s equations coupled with a plasma response from the fluid motion of electrons), but it uses a Fourier method in the plane normal to the density gradient. A more detailed account of the other codes can be found in [54].

All the codes are cold-plasma codes, encapsulating the same physics, it is just their implementation that differs. The main difference between FFW and the two FDTD codes is that the FFW code launches a Gaussian beam all with the exact same wave vectors (so truly pure O-mode) whereas the two FDTD codes launch a Gaussian beam whose centre is pure O-mode, but with a range of wave vectors throughout the beam such that the tails of the beam are no longer pure O-mode.

In order to compare the codes, a series of benchmark scenarios of increasing complexity were decided on. A summary of these can be found in Table 5.1.

A summary of the numerical parameters used in each scenario for EMIT-2D can be found in Table 5.2. The parameters used in the other codes were different, decided by my collaborators without consultation between us as the primary purpose of this benchmark was to determine if we would arrive at the same results each using our own processes.

Scenario	Resolution	N_x	N_z	CFL	N_t	ν/ω_0
1	50	4000	1500	0.25	20000	0
2	50	4000	1500	0.25	20000	0
3	50	4000	1500	0.25	20000	0
4	50	4000	1500	0.25	20000	0
5	50	4000	1500	0.25	20000	1×10^{-4}
6	50	4000	1500	0.25	20000	1×10^{-4}

Table 5.2: A summary of the numerical parameters used in the benchmark scenarios by EMIT-2D. Resolution is the number of Yee cells per vacuum wavelength, and N_x and N_z are the number of Yee cells over which the simulation domain extends. CFL relates the time step to the grid spacing, and N_t is the total number of timesteps over which the simulation was run. Finally, ν is the numerical collision frequency normalised to the vacuum wave frequency.

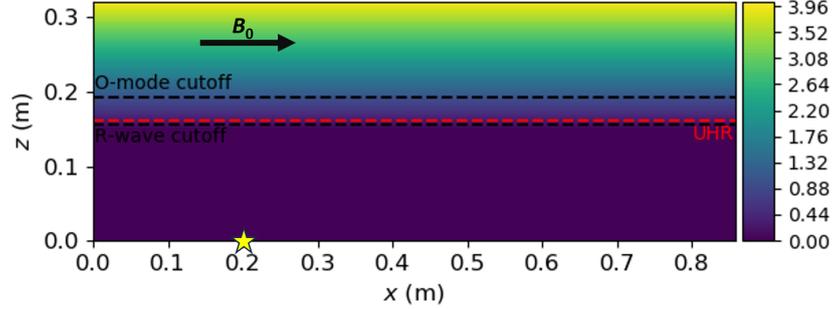


Figure 5.1: Simulation domain for scenarios 4, 5 and 6. The colourmap indicates normalised density $X = n_e/n_{\text{crit}}$. An arrow indicates the direction of the background field, $B_0 = 0.85$ T corresponding to $Y \approx 0.5$, and a star indicates the position of the antenna.

In cases of perpendicular propagation, the beam propagates purely in the z -direction. Where an angle of propagation is specified, it is the angle to the z -axis. In the case where a plasma is present, the beam is still launched in vacuum, with a linear density gradient starting mid-way into the simulation domain with a normalised density scale length of $k_0 L_n = 25$. This density profile can be described as

$$n_e(z > z_{\text{start}}) = (z - z_{\text{start}}) \frac{2\pi}{\lambda_0} \frac{1}{k_0 L_n} = \frac{z - z_{\text{start}}}{L_n} \quad (5.1)$$

where $z_{\text{start}} = 0.15$ m is the location in z at which the plasma starts, and all other symbols have their previously given meanings. A plot of this density profile can be seen in Fig. 5.1 with the antenna location and B -field direction (when present) also shown.

In the cases where a background magnetic field is present (scenarios 5 and 6) it was set to a uniform value of $B_0 = 0.85$ T (or $Y = \frac{\omega_{ce}}{\omega} \approx 0.5$) pointing in the x -direction throughout

the simulation domain.

In all cases, the launched beam has a frequency of $f_0 = 28$ GHz, propagating either entirely in the z -direction or at an acute angle to the z -axis. In the case where there is a magnetised plasma present, an O-mode beam is launched using the new elliptical polarisation functionality in EMIT-2D to match the elliptical polarisation used by the other codes. The presence of a magnetised plasma also introduces cut-offs and resonances. For the purposes of the OX-mode conversion scheme, the main cut-off of interest is the O-mode cut-off, however it is important to note that the R-wave cut-off exists as well at a slightly lower density. Between these two cut-offs is the UHR where, in a warm plasma code, the converted X-mode would convert to an EBW. However, in a cold plasma code this is not possible. Instead, to dissipate the power that would build up there, a numerical collisionality is used. This is the same method used to damp the waves in the boundaries of the simulation domain, however, the value used for the parameter within the simulation domain is significantly smaller, and is set to a constant value throughout the domain. The same method is used by IPF-FDMC [58], while FFW instead resolves the small scale electrostatic waves generated at the UHR which carry the power away.

The comparison between the three codes was carried out by looking at the power of the microwave beam at various slices in z .

5.1.2 Scenario 1

The first scenario considered was a diverging Gaussian beam, propagating through vacuum purely in the z -direction with its focus at the antenna where it has a beam waist of $4\lambda_0$. The power in the electric field at various slices in z can be seen in Fig. (5.2), along with the difference between the codes at these locations.

In Fig. (5.2a), we can see the peak amplitude of the beam decreasing with z as it diverges, spreading out its power over a larger area, as expected. In Fig. (5.2b) we can see that the agreement is best between the two FDTD codes, EMIT-2D and IPF-FDMC, with the difference between the codes remaining below 0.7% of the measured power. This is to be expected as they are the most similar, employing different implementations of the same algorithm. The agreement between the two FDTD codes and the FFW code is also good, remaining below 5% of the measured power.

It is possible that improving the resolution in all three codes (increasing the number of

grid points per wavelength) would result in even better agreement, but this would greatly increase the computational cost. In addition, whilst comparing the power in the electric field at slices in z is a good initial benchmark, we are more concerned with whether the end result parameters (such as OX mode conversion efficiency in Scenario 5 and 6) agree using three different codes where the users have chosen numerical parameters, such as resolution and time-step, independently of each other.

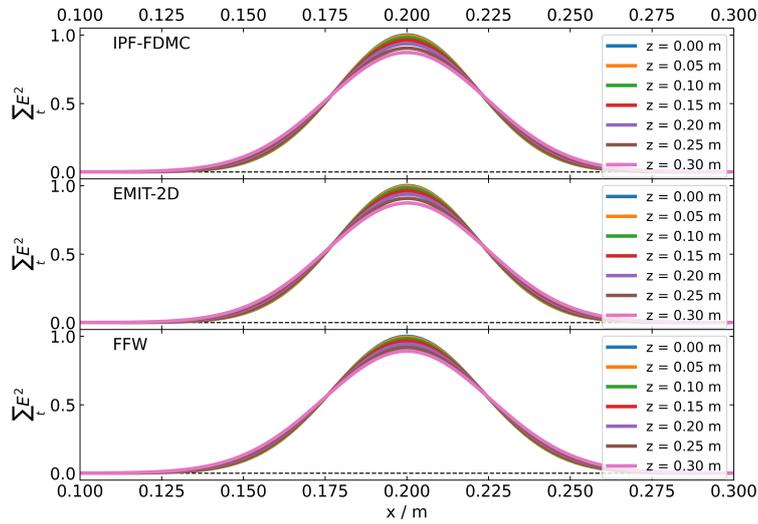
As a further check, the total power in the electric field in each slice was calculated for EMIT-2D, which we would expect to be conserved in vacuum. The percentage change in this value can be seen in Fig. 5.3. As expected, the value is approximately conserved, with no change exceeding a magnitude of 0.3 % of the original value in the antenna plane. The point corresponding to the z -slice within the damping boundaries of the code is omitted from the plot, but has a percentage change of -99.8 % as would be expected.

5.1.3 Scenario 2

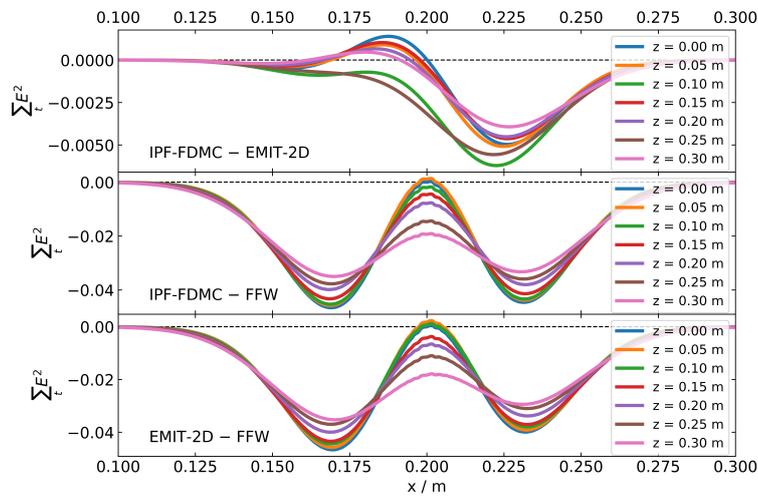
The next scenario considered was a converging Gaussian beam, propagating through vacuum purely in the z -direction with its focus 0.2 m into the simulation domain where it has a beam waist of $4\lambda_0$. The power in the electric field at various slices in z can be seen in Fig. (5.4), along with the difference between the codes at these locations.

In Fig. (5.4a), we can see the peak amplitude of the beam increasing with z as it converges, getting narrower as expected. In Fig. (5.4b) we can see that just like for scenario 1, the agreement is best between the two FDTD codes, EMIT-2D and IPF-FDMC, with the difference between the codes remaining below 0.7% of the measured power. The agreement between the two FDTD codes and the FFW code is also good, remaining below 1% of the measured power.

As with scenario 1, the total power in the electric field in each slice was calculated for EMIT-2D. The percentage change in this value can be seen in Fig. 5.3. As expected, the value is approximately conserved, with no change exceeding a magnitude of 0.31 % of the original value in the antenna plane. The point corresponding to the z -slice within the damping boundaries of the code is omitted from the plot, but has a percentage change of -99.8 % as would be expected.



(a) Microwave beam power at various slices through the simulation domain



(b) Difference in the power signal between the codes.

Figure 5.2: Full-wave simulation results for scenario 1. A Gaussian beam is launched with its focus at the antenna, propagating in the z -direction through a vacuum. The power in the beam is read out at different slices in z and compared between the codes.

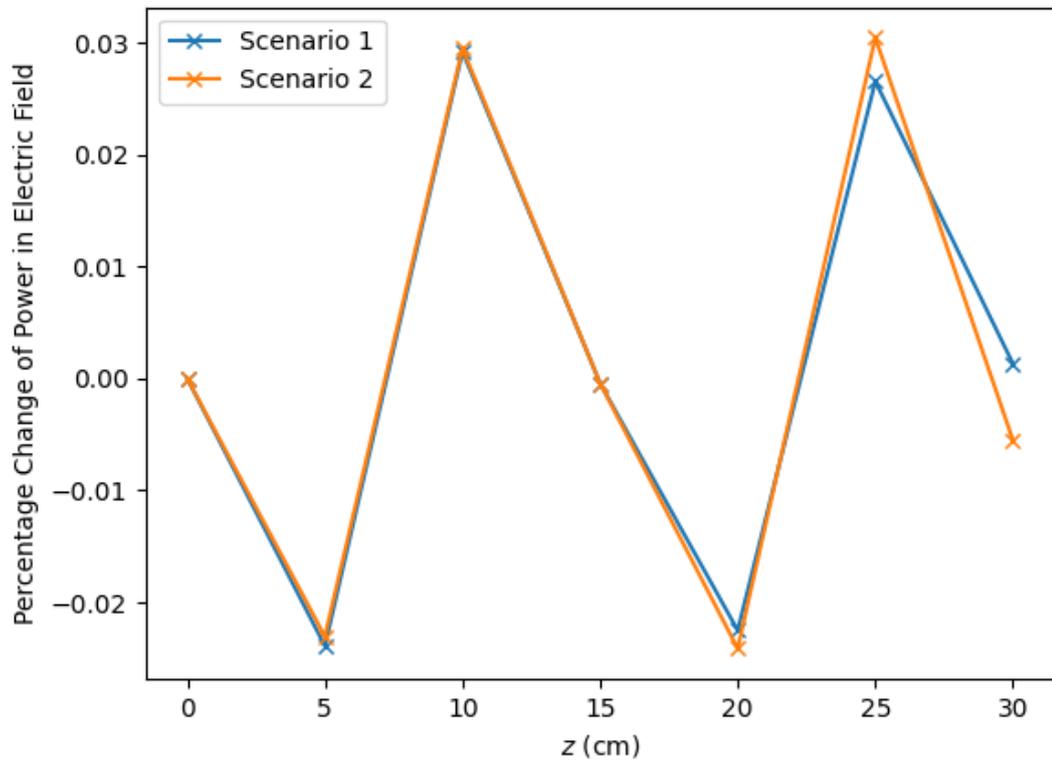
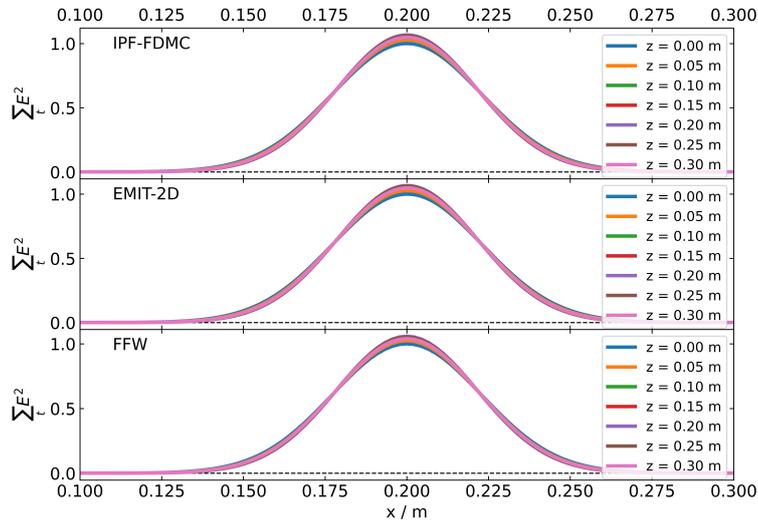
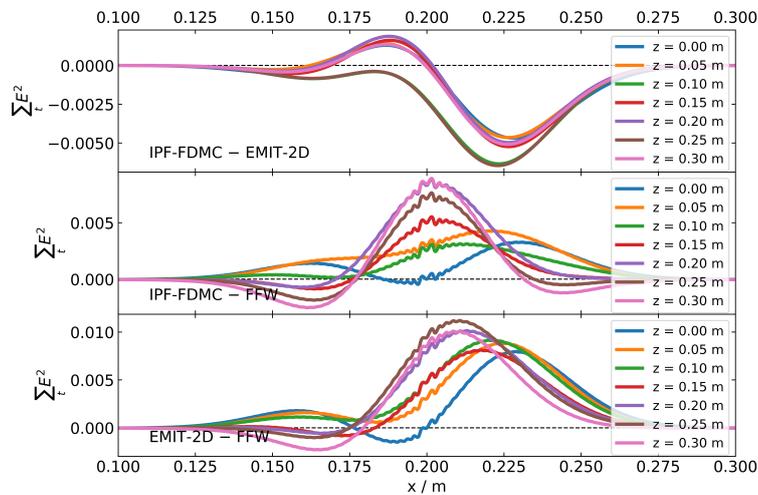


Figure 5.3: Percentage change in the total power in the electric field at different slices in z for EMIT-2D scenarios 1 and 2. The point corresponding to the z -slice within the damping boundaries is omitted to avoid skewing the axis of the plot. The percentage change outside of the damping boundaries does not exceed a magnitude of 0.31 % of the original value in the antenna plane for any of the scenarios.



(a) Microwave beam power at various slices through the simulation domain



(b) Difference in the power signal between the codes.

Figure 5.4: Full-wave simulation results for scenario 2. A Gaussian beam is launched with its focus 0.2 m into the plasma, propagating in the z -direction through a vacuum. The power in the beam is read out at different slices in z and compared between the codes.

5.1.4 Scenario 3

The next scenario considered reverted to a diverging Gaussian beam with a beam waist of $4\lambda_0$ at the antenna. The beam was launched to propagate through vacuum at an angle of 30° to the z -axis. The power in the electric field at various slices in z can be seen in Fig. (5.5), along with the difference between the codes at these locations.

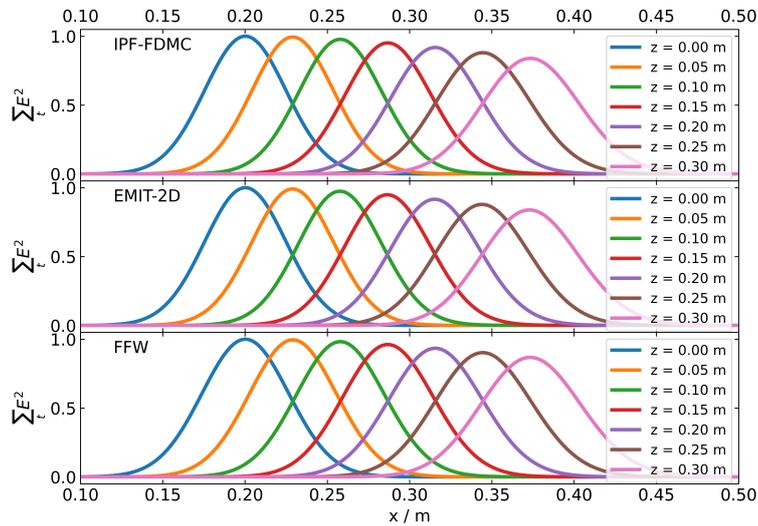
In Fig. (5.5a), we can see the peak amplitude of the beam decreasing with z as it diverges, while it also shifts to the left as expected for a beam propagating at an angle. In Fig. (5.5b) we can see that just like for scenarios 1 and 2, the agreement is best between the two FDTD codes, EMIT-2D and IPF-FDMC, with the difference between the codes remaining below 1% of the measured power. The agreement between the two FDTD codes and the FFW code is also good, remaining below 5% of the measured power.

We no longer expect the total power in the electric field at each z -slice to be conserved, as the beam is now propagating at an angle to z , hence it was not calculated.

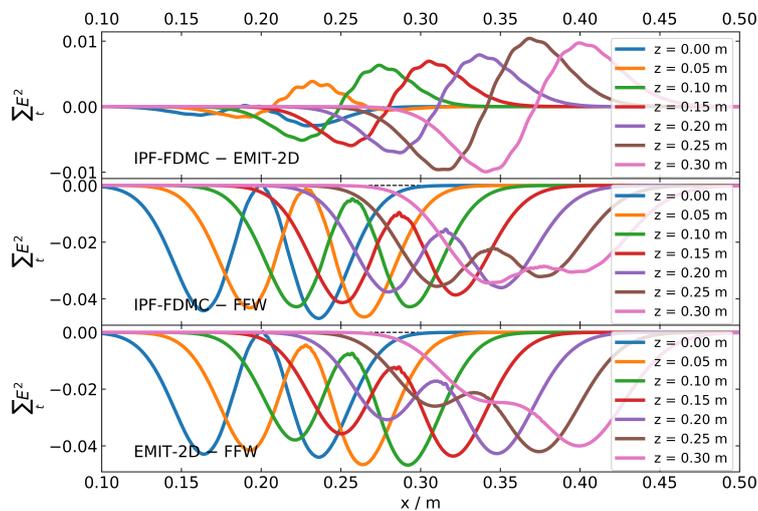
5.1.5 Scenario 4

After studying the previous three vacuum cases, we then moved on to the scenarios where a plasma was present, as shown in the density profile shown in Fig. (5.1). In the case of this scenario, there was still no background magnetic field. In this set-up, we launched a diverging Gaussian beam with a beam waist of $4\lambda_0$ at the antenna, propagating at an angle of 30° to the z -axis. The power in the electric field at various slices in z can be seen in Fig. (5.6), along with the difference between the codes at these locations.

In Fig. (5.6a), we can now see two peaks at each z -location — the injected and the reflected beam. The peak amplitude of the reflected beam is clearly lower than the amplitude of the injected beam in each case, as expected. We can also see that there is no longer any signal beyond $z \geq 0.2$ m as the beam is reflected at the cut-off density before this. In Fig. (5.6b) we can see that the differences between the codes are most pronounced close to the cut-off. Initially, we supposed that this could be due to a misalignment of the numerical grids used by the various codes, placing the cut-off at slightly different locations relative to the z -location at which the field is read. When very close to the cut-off, the shape of the signal with respect to x depends very strongly on just how close to the cut-off it is and the position at which it hits the cut-off, due to the fact that interference is taking place between the incident and reflected beam, so moving the detected signal location slightly closer to the

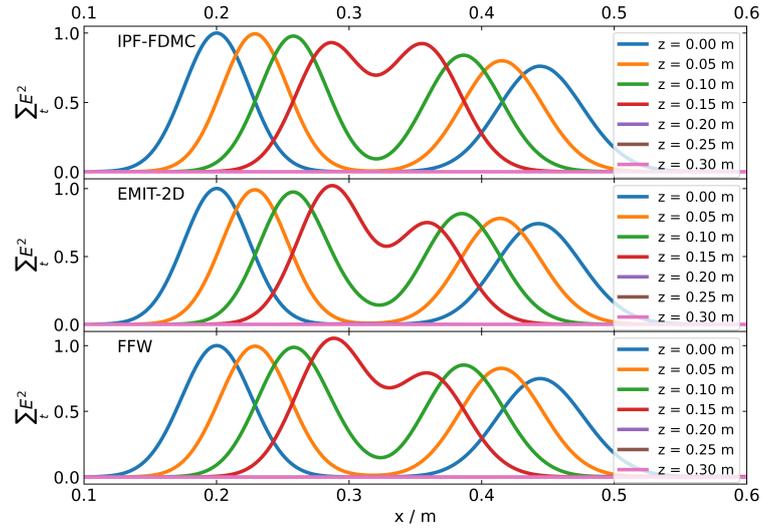


(a) Microwave beam power at various slices through the simulation domain

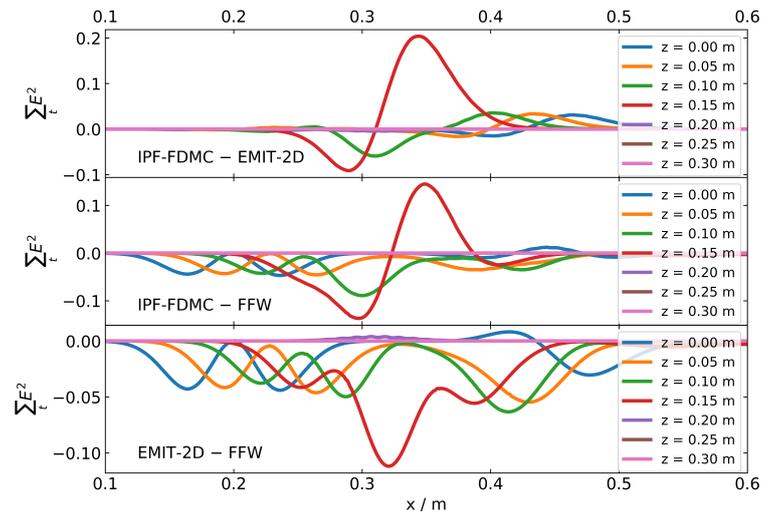


(b) Difference in the power signal between the codes.

Figure 5.5: Full-wave simulation results for scenario 3. A Gaussian beam is launched with its focus at the antenna, propagating at 30° to the z -axis through a vacuum. The power in the beam is read out at different slices in z and compared between the codes



(a) Microwave beam power at various slices through the simulation domain



(b) Difference in the power signal between the codes.

Figure 5.6: Full-wave simulation results for scenario 4. A Gaussian beam is launched in vacuum with its focus at the antenna, propagating at 30° to the z -axis into a plasma with a linear density gradient. It is reflected at the cut-off density in the plasma before propagating back to the antenna plane. The power in the beam is read out at different slices in z and compared between the codes

cut-off (by as little as 1 mm for example) was expected to have a noticeable effect.

In order to test whether the difference was due to misalignment of numerical grids caused by differing resolutions, I reran the simulation in EMIT-2D using different resolutions, however no great change was observed. For example, doubling the resolution, as shown in Fig. (5.7), causes a maximal difference of order 1% in the slice closest to the cut-off, significantly

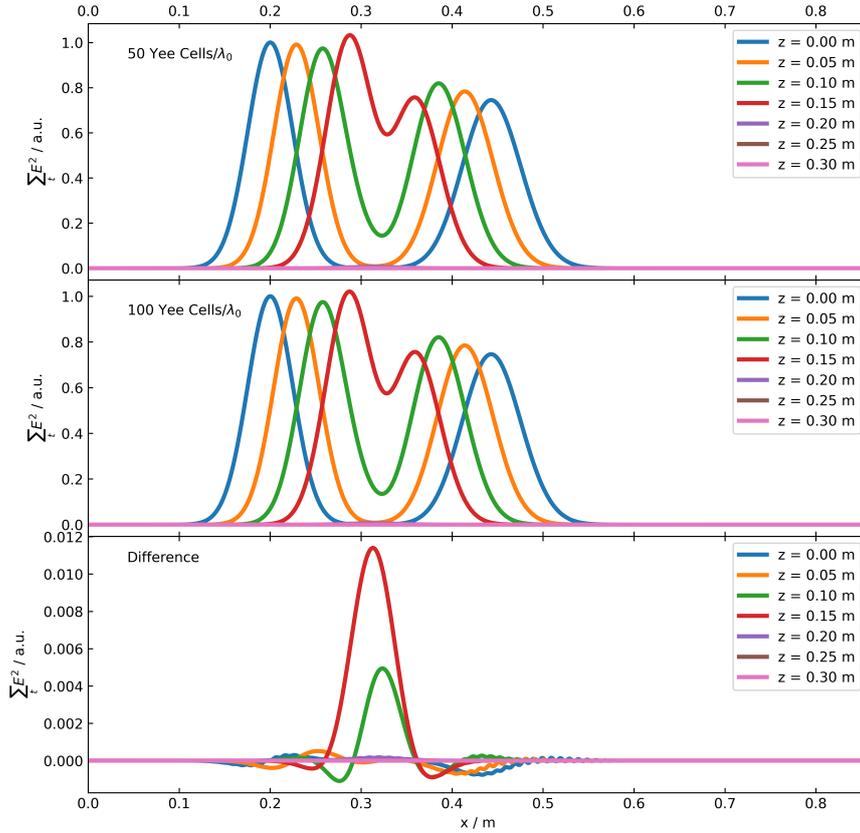


Figure 5.7: Full-wave simulation results for scenario 4 using EMIT-2D with two different resolutions. A Gaussian beam is launched in vacuum with its focus at the antenna, propagating at 30° to the z -axis into a plasma with a linear density gradient. It is reflected at the cut-off density in the plasma before propagating back to the antenna plane. The power in the beam is read out at different slices in z and compared between the resolutions. The top figure shows microwave beam power at various slices through the simulation domain when using a resolution of 50 Yee Cells per wavelength. The middle figure shows the microwave beam power at various slices through the simulation domain when using a resolution of 100 Yee Cells per wavelength. The bottom figure shows the difference between the two at various slices through the simulation domain. This difference never exceeds 1.2% and is highest in the slice closest to the cut-off, where we expect constructive and destructive interference to be occurring, making the code more sensitive to slight changes in relative position between the cut-off surface and the location we are reading the field due to the discrete grid.

less than the of order 10% difference observed between the different codes.

Instead, the discrepancy in codes was determined to be due to the polarisation of the beam. As there is no magnetic field, we do not expect the dispersion relation to depend on beam polarisation, provided we are not close to a cut-off. However, near a cut-off, the direction of the polarisation vector with respect to direction of propagation and density gradient will be important. In order to determine this, I reran the simulation in EMIT-2D

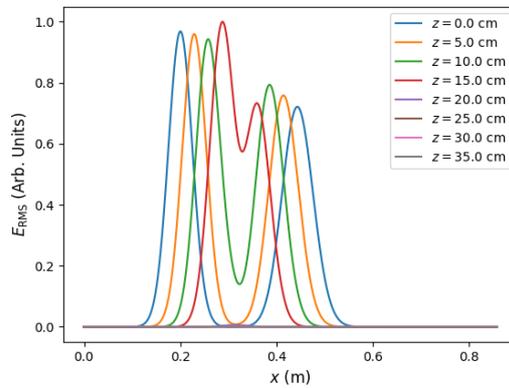
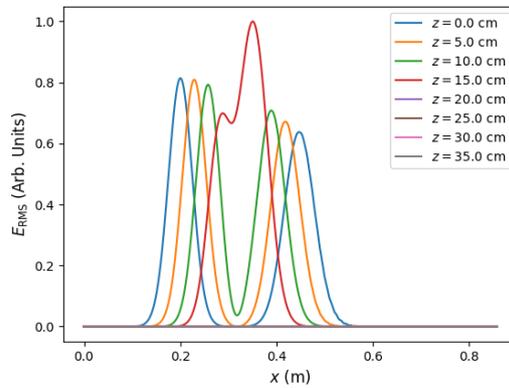
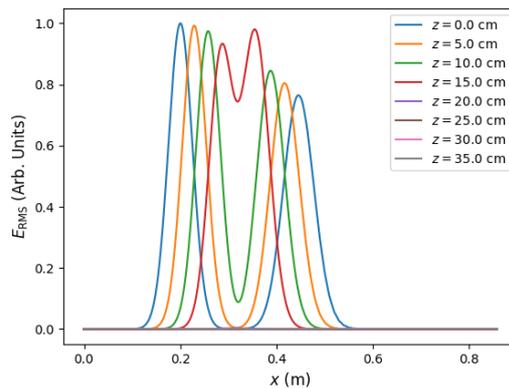
using three different polarisations: the original x -polarised beam, an elliptically polarised beam, and a y -polarised beam. The results of these simulations can be seen in Fig. (5.8). As can be seen in the figure, changing the polarisation has minimal effect on power in the electric field read out at different slices in z apart from when near the cut-off. Therefore, this is likely the cause of the disagreement between the codes.

IPF-FDMC automatically uses an elliptical polarisation when launching at an angle. Comparing the three codes utilising an elliptical polarisation in EMIT-2D as well gives the results shown in Fig. (5.9). As can be seen, the large discrepancy between the two FDTD codes and FFW remains, as now FFW is the only code with a linear polarisation, but the difference between EMIT-2D and IPF-FDMC has reduced to have a maximum of around 5%. In future simulations for scenarios 5 and 6 an elliptical polarisation was used for all three codes.

5.1.6 Scenario 5

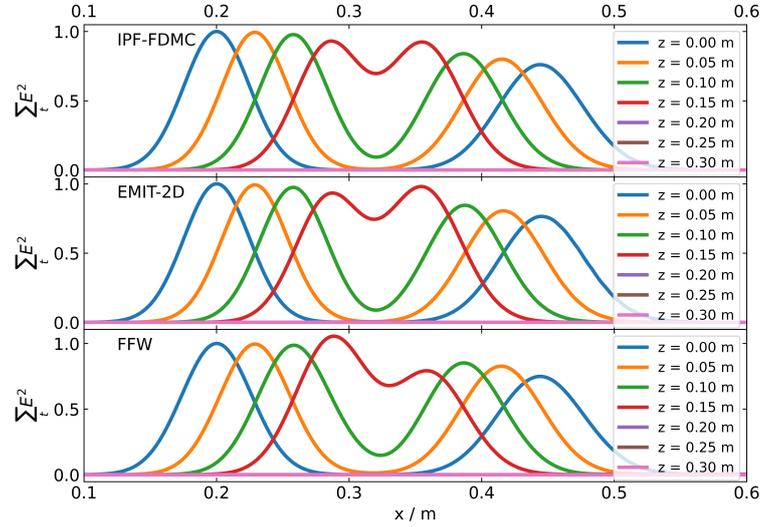
This scenario still used the density profile shown in Fig. (5.1), while also adding a background field of 0.85 T. The launch conditions were to use a diverging beam propagating at 30° to the z -axis, however, now that there is a magnetic field present, in order to launch pure O-mode, the elliptical polarisation must be used (rather than linear polarisation in the x -direction). The power in the electric field at various slices in z can be seen in Fig. (5.10), along with the difference between the codes at these locations.

In Fig. (5.10a), we can see two peaks at each z -location — the injected and the reflected beam. The peak amplitude of the reflected beam is clearly lower than the amplitude of the injected beam in each case, as expected. We can also see that there is no longer any signal beyond $z \geq 0.2$ m as the beam is reflected at the cut-off density before this. The peak amplitude across the slices is in the slice closest to the cut-off. This is due to a power increase at the UHR which must be dissipated using numerical collisionality. In Fig. (5.10b) we can see that as in previous scenarios, agreement is best between EMIT-2D and IPF-FDMC as expected as they are both different implementations of the same algorithm. The difference between these two codes and FFW is most pronounced between the O-mode cut-off and the UHR. We do not necessarily expect the codes to agree here, as they are not accurately modelling the physical processes that would be taking place, as this would require a warm-plasma code. The exact value of the electric field in this location is likely to be sensitive to

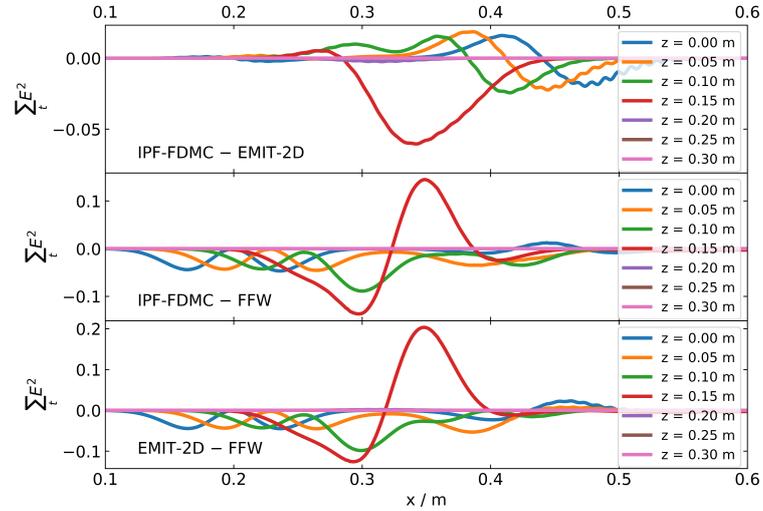
(a) Linearly polarised in the x -direction(b) Linearly polarised in the y -direction

(c) Elliptically polarised

Figure 5.8: Full-wave simulation results for scenario 4 using EMIT-2D with three different polarisations. A Gaussian beam is launched in vacuum with its focus at the antenna, propagating at 30° to the z -axis into a plasma with a linear density gradient. It is reflected at the cut-off density in the plasma before propagating back to the antenna plane. The power in the beam is read out at different slices in z .



(a) Microwave beam power at various slices through the simulation domain



(b) Difference in the power signal between the codes.

Figure 5.9: Full-wave simulation results for scenario 4, now using an elliptical polarisation in EMIT-2D. A Gaussian beam is launched in vacuum with its focus at the antenna, propagating at 30° to the z -axis into a plasma with a linear density gradient. It is reflected at the cut-off density in the plasma before propagating back to the antenna plane. The power in the beam is read out at different slices in z and compared between the codes.

the exact placement of the cut-off and resonance relative to the slice in z at which we are reading the field out. It is also likely to be sensitive to the value chosen for the artificial collisional frequency used by EMIT-2D and IPF-FMDC used to dissipate the power that would otherwise build up at the UHR as well as being sensitive to the exact form of the algorithm used to propagate the wave.

The OX mode conversion efficiency could then be calculated by comparing injected power to reflected power in the antenna plane. Any missing power is then assumed to have been successfully converted to X-mode. EMIT-2D gives an efficiency of 1.9 % while IPF-FMDC gives 1.1 % and FFW gives 0.2 %. These values are low as we are far from the optimum angle.

5.1.7 Scenario 6

This scenario still used the density profile shown in Fig. (5.1), with a background field of 0.85 T, just like in scenario 5. However, for this scenario we injected at the optimum angle for OX mode conversion efficiency, as given by [52]

$$\varphi_{\text{opt}} = \arccos \left[\left(\frac{Y}{1+Y} \right)^{1/2} \right] \quad (5.2)$$

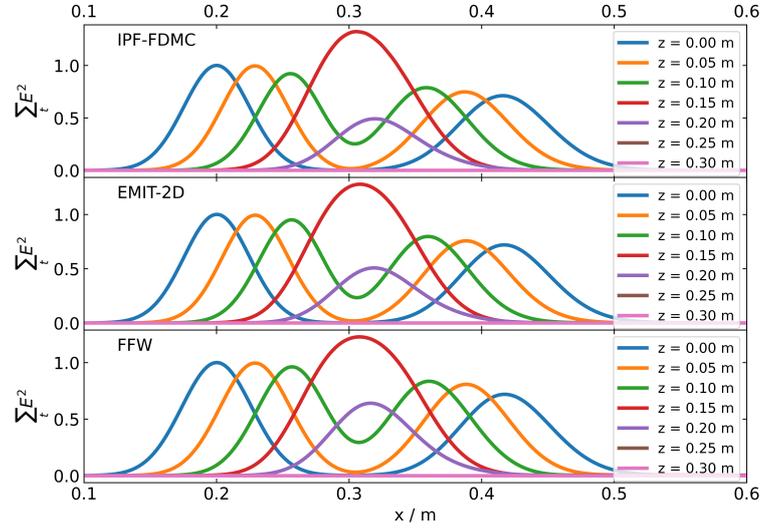
where this angle is defined in reference to the magnetic field. Therefore, in the set-up described where B_0 points in the x -direction, our optimum angle of injection to z will be

$$\theta_{\text{opt}} = 90^\circ - \varphi_{\text{opt}} \quad (5.3)$$

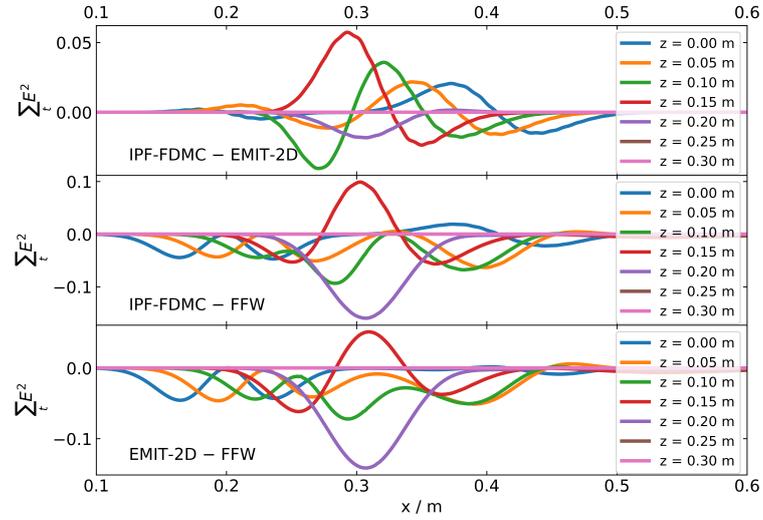
in degrees. For a frequency of 28 GHz and a background field of 0.85 T, this equates to $\theta_{\text{opt}} \approx 47^\circ$.

The power in the electric field at various slices in z can be seen in Fig. (5.11), along with the difference between the codes at these locations.

In Fig. (5.11a), we can see that the amplitude of the reflected peak is now greatly reduced due to significant OX mode conversion taking place. This results in significantly less O-mode power being reflected back out of the plasma. The peak amplitude across the slices is in the slice closest to the cut-off. This is due to a power increase at the UHR which must be dissipated using numerical collisionality. Another important feature is the ‘hole’ in the reflected (non-converted) part of the beam which can also be seen in Fig. (5.12). This is



(a) Microwave beam power at various slices through the simulation domain



(b) Difference in the power signal between the codes.

Figure 5.10: Full-wave simulation results for scenario 5. A diverging O-mode Gaussian beam is launched in vacuum with uniform background magnetic field in the x -direction, propagating at 30° to the z -axis into a plasma with a linear density gradient. At the O-mode cut-off density, the O-mode is reflected but some mode conversion to the X-mode also occurs. The power in the beam is read out at different slices in z and compared between the codes.

due to the converted SX-mode being filtered out by the UHR (where it would convert to an EBW in a warm plasma code) where it is dissipated via the numerical collisional damping term in the codes.

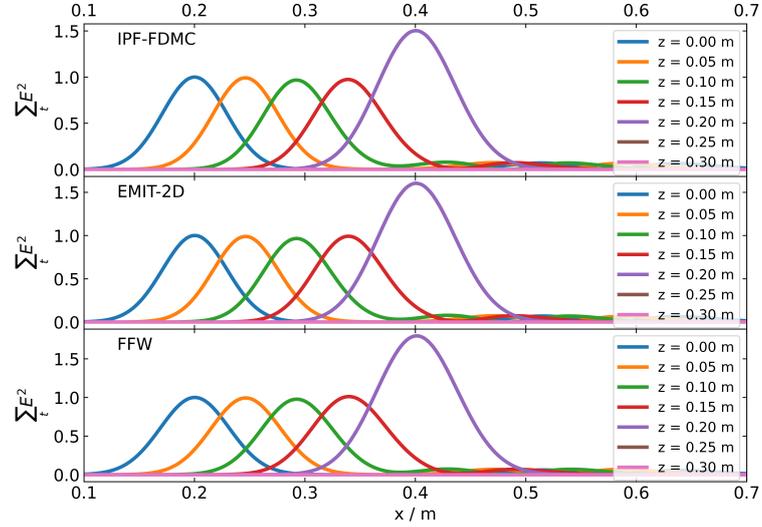
In Fig. (5.11b) we can see that as in previous scenarios, agreement is best between EMIT-2D and IPF-FDMC as expected as they are both different implementations of the same algorithm. The difference between these two codes and FFW is most pronounced between the O-mode cut-off and the UHR, as was the case in scenario 5, and the reasoning as to why is believed to be the same here. As the codes are not actually modelling the physics that would be occurring here, we are not particularly concerned with the exact values of the field here. The main parameter of interest is the OX-mode conversion efficiency.

The OX mode conversion efficiency was calculated by comparing injected power to reflected power in the antenna plane. Any missing power is then assumed to have been successfully converted to X-mode. Now that we are at the optimum angle, we expect much higher conversion efficiencies than in scenario 5. EMIT-2D gives an efficiency of 86.7 % while IPF-FDMC gives 86.4 % and FFW gives 88 %. The higher value obtained from the FFW code is due to its implementation of a Gaussian beam using a spectrum of k -components, each of which is in pure O-mode polarisation. The two FDTD codes implement the Gaussian beam via spatially varying wave electric field where the centre of the beam will be pure O-mode but the tails will be in non-optimum polarisation (as one might expect from a physical antenna). The differences between the values obtained from all three codes can be used to estimate the numerical error on the OX mode conversion efficiency at approximately $\pm 2\%$.

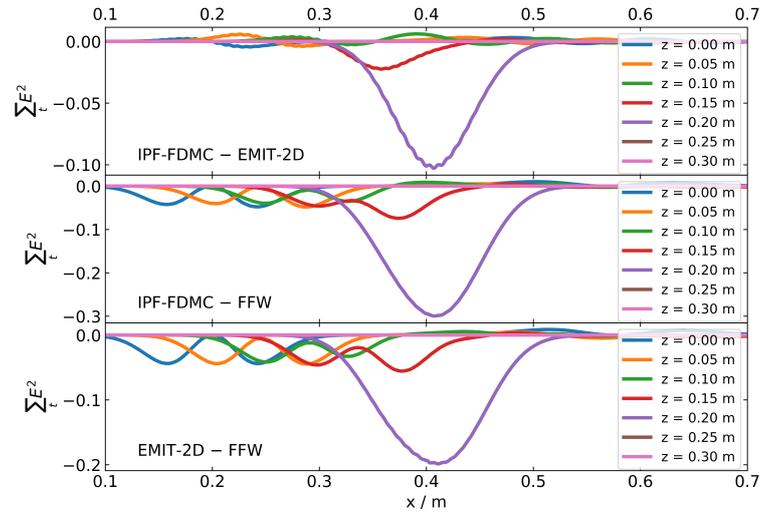
Overall, the codes are deemed to agree well across the range of scenarios, apart from when proximity to the cut-off makes the exact details of polarisation important, and when considering the region between the O-mode cut-off and the UHR, in which we do not expect the output of the codes to be wholly physical anyway.

5.2 Effect of Density Scale Length on OX Mode Conversion

As mentioned previously, part of the motivation for this work was a project to install gyrotrons on MAST-U, developing a suite of codes that could be applied to optimisation studies and comparison to experiments. The process of optimising the current drive achievable by the microwave system required considering different antenna positions for a possible 28



(a) Microwave beam power at various slices through the simulation domain



(b) Difference in the power signal between the codes.

Figure 5.11: Full-wave simulation results for scenario 6. A diverging O-mode Gaussian beam is launched in vacuum with uniform background magnetic field in the x -direction, propagating at the optimum angle for OX mode conversion efficiency, which equates to an angle of 47° to the z -axis into a plasma with a linear density gradient. At the O-mode cut-off density, mode conversion to the X-mode occurs, and any O-mode that is not converted is reflected back out of the plasma. The power in the beam is read out at different slices in z and compared between the codes

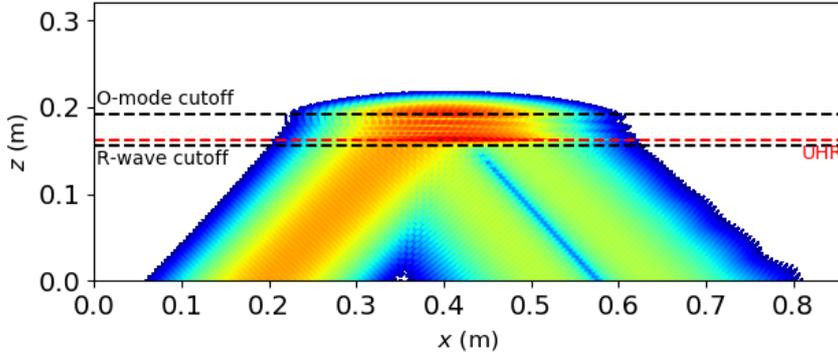


Figure 5.12: Snapshot of the wave power for scenario 6. Colour represents power in arbitrary units on a logarithmic scale. The diverging Gaussian beam can be seen travelling toward the O-mode cut-off at the optimum angle for mode conversion. At the O-mode cut-off, the majority of the beam is converted to the SX-mode. In a warm plasma, this would then travel to the UHR where it converted to an EBW before propagating further into the plasma. but, in this cold plasma code, the converted X-mode is instead dissipated with numerical collisionality. The tails of the injected O-mode beam do not have the optimum wave-vector so are not converted to the SX-mode. Any unconverted O-mode is reflected back, resulting in the hollow middle of the reflected beam.

GHz launcher to heat both L-mode and H-mode plasmas using EBWs. The presence of the pedestal in H-mode plasmas means that there will potentially be very steep density gradients in the edge where the OX mode conversion is taking place.

To this end, as a final stage of the code comparison, we considered varying density scale lengths, $k_0 L_n$, for the linear density profile. The simulation set-up was the same as for scenario 6, injecting at the optimum angle and calculating the mode conversion efficiency in each case. Based on high elongation, low plasma- β beta scenarios in MAST-U, the range decided on was $k_0 L_n = 2 - 25$. The results of this scan can be seen in Fig. 5.13. As before, excellent agreement is seen between the three codes.

It is expected that decreased $k_0 L_n$ will increase mode conversion efficiency, as steeper density gradients are less sensitive to any slight angular mismatch which results from the divergence of the microwave beam as it travels to the mode conversion layer. This is because

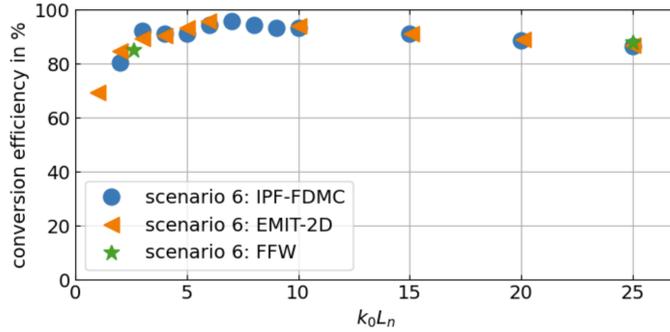


Figure 5.13: How OX mode conversion efficiency depends on density scale length of a linear density profile in the presence of uniform background magnetic field.

a steeper density gradient means that the evanescent layer the O-mode will have to tunnel through to reach the OX mode conversion point will be thinner. Our results reproduced this in the range $k_0 L_n = 5 - 25$ where a slow decrease in conversion efficiency can be seen with increasing $k_0 L_n$.

The sharp decrease for steeper density profiles $k_0 L_n < 5$ was initially surprising. The case for $k_0 L_n = 3$ is shown in more detail in Fig. 5.14. Looking at the electric field in the antenna plane for scenario 6, a clear dip is seen in the beam centre where the missing power has been converted into an X-mode wave. This dip is absent for the shorter density scale length. We determined that this was due to the successfully mode-converted SX-mode being able to tunnel through the evanescent layer between the UHR and the R cut-off and couple to the FX-mode which then propagates out of the plasma. This is only possible when a steep density gradient decreases the width of the evanescent layer sufficiently such that it is of comparable length to the wavelength of the wave to allow significant tunnelling.

5.3 Summary and Conclusions

In summary, a comparison of three different full-wave codes EMIT-2D, IPF-FDMC, and FFW has yielded good agreement in a range of scenarios with increasing complexity. This gives me good confidence in the newly developed EMIT-2D for use in the research projects described in Chapters 6 and 7.

As part of the benchmark, a brief study into the effect of density scale length on OX mode conversion efficiency was carried out over parameter ranges relevant to MAST-U. This reproduced the expected result that OX mode conversion efficiency decreased for shallower density

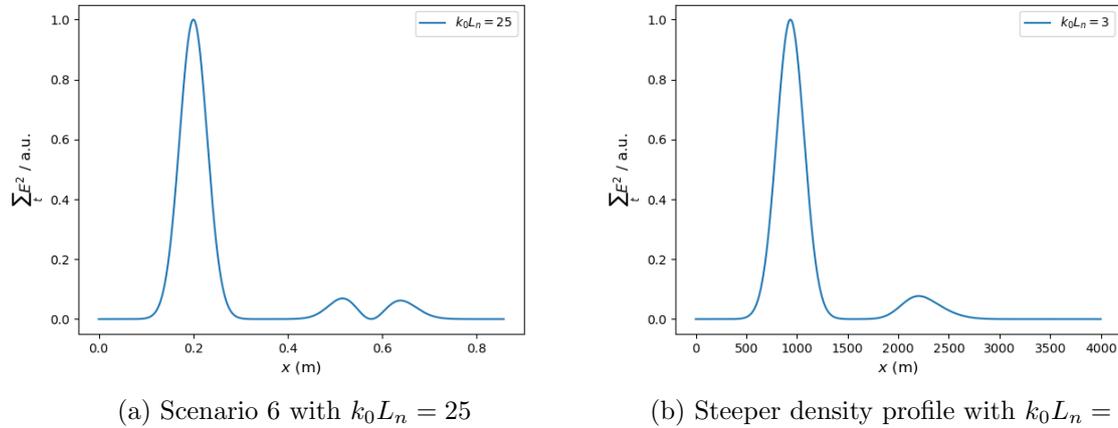


Figure 5.14: Power in the antenna plane for two scenarios with different density scale lengths. The large peak to the left in each plot is the emitting antenna. The smaller peaks to the right are the reflections, indicating the parts of the beam that were not successfully mode converted. The gap in the middle of the reflected beam in figure (a) is typical for a case where OX mode conversion has taken place. This is absent from figure (b) due to the X-mode tunnelling back out of the plasma along with the reflected O-mode.

gradients and revealed that it also decreased for very steep profiles where the converted X-mode was able to tunnel through the evanescent layer. This is a potentially important point for the MAST-U gyrotron project when considering the path of reflected power and the damage it could cause to in-vessel components.

Chapter 6

ECRH Beam Broadening on DIII-D

This work was the subject of a paper published at M.W. Brookman, *et al 2023 Nucl. Fusion* **63** 044001 DOI: 10.1088/1741-4326/acbb8e [59]. As such, much of the following chapter is reproduced directly with the permission of Nuclear Fusion and EUROfusion.

When I started work on my PhD, this project was already in progress, with simulations having been completed by a previous PhD student using the version of EMIT-3D I inherited at the start of my PhD. However, the referees raised some valid concerns with some of the assumptions that had gone into the simulations. These assumptions had been made to reduce computational cost while using EMIT-3D. As such, when I started work on the project, my role was to come up with a new methodology for the simulations which would not rely on these assumptions. This was the main motivation for developing EMIT-2D.

The experimental measurements of beam broadening and the measurements of plasma parameters such as background density, fluctuation amplitude, and turbulence length scales were undertaken by collaborators. I then used the experimental information I had about the plasma to create turbulence profiles and run simulations through them using EMIT-2D to predict the beam broadening effect of the turbulence.

6.1 Motivation

As discussed in Sec. 1.3, microwaves have many uses in tokamaks. They are used in diagnostics such as ECE to measure the plasma temperature, and they are used for Doppler

back-scattering reflectometry to characterise plasma turbulence [60]. They are also used for heating and current drive where a high-power beam is injected with a frequency chosen to match a particular resonance at the targeted absorption location [61,62].

The next generation of tokamaks, including ITER [63], will rely on high-power microwave injection for global heating and for driving local currents to control and mitigate magneto-hydrodynamic (MHD) instabilities. Local microwave power injection will be key in stabilising neoclassical tearing modes (NTMs), magnetic reconnection instabilities which arise from the interplay between local reductions in plasma current and a flattening of the pressure profile, resulting in a degradation of confinement of the device. They can be stabilised by driving a small local current to counteract the local reduction in plasma current, but this requires precise spatial localisation of the current driven by the microwave beam [17,64–66].

However, as discussed in Sec. 1.2, steep temperature and density gradients near the edge of fusion devices drive a number of instabilities, resulting in a layer of turbulence where the density fluctuation level can exceed 100% of the background density [12]. These density fluctuations are often on comparable length scales to microwave wavelengths. This can cause significant scattering of incident microwaves and an overall broadening of beams.

This has been a topic of interest since high-power microwave sources became available for plasma heating in the 1980s. Initially, numerical studies were conducted using ray tracing codes in order to predict the scattering effect density fluctuations may have on an injected microwave [67–69]. It was found that even for fluctuation levels as low as 1% significant scattering could be expected but was determined to only be a critical issue in cases where localised deposition was of importance, such as for the stabilisation of NTMs.

The topic regained focus after ITER-like scenarios of beam broadening by edge turbulence were considered by Tsironis *et. al* (2009) [70] and it was found that it was possible the additional broadening from the plasma turbulence could result in a doubling of beam width in ITER-like scenarios [70]. Such a broadening could significantly impact the efficiency of the microwave heating and current drive systems by depositing power where it will be absorbed less efficiently, and wasting power building an undesired current drive profile, and consequentially could significantly lessen the efficiency of the whole device. This will be a key consideration when attempting to achieve net generation of power.

Furthermore, in the case of NTM stabilisation, broadening or misaligning of the microwave beams used leads to wasted power, reducing the machine efficiency, and may also

lead to additional requirements on the microwave systems such as additional modulation of the microwave power in real-time [71]. In the worst-case scenario, it could also lead to a failure to stabilise the NTM, leading to a critical loss of plasma confinement.

The prediction of the broadening of microwave beams by plasma turbulence is therefore of great importance for the design and operation of future tokamaks. Studying this beam broadening experimentally and with simulations is key to improve our understanding of the effect and is an active area of research [28, 53, 70–76]. To this end, we have carried out an investigation comparing experimentally measured beam broadening by plasma turbulence to full-wave simulations using EMIT-2D. This benchmark of simulation against experiment will lead to greater confidence in predictions for future devices, and is the first time such a study has been carried out.

6.2 Experimental Measurements

A more detailed description of the experimental work can be found in [77]. It was conducted by collaborators at DIII-D rather than by myself but is included here for completeness.

This work is based on experimental observations on the DIII-D tokamak plasma, with typical parameters of [78]:

- major radius, $R = 1.7$ m
- minor radius, $a = 0.67$ m
- plasma current, $I_p = 800$ kA
- on-axis toroidal magnetic field, $B_t = 2$ T

The gyrotron beams are launched from a set of steerable mirrors located 60° above the midplane on the tokamak's outboard side, as shown in Fig. 6.1, and are optimised for X-mode 2nd harmonic EC absorption at a magnetic field strength of 2 T with the following parameters [79, 80]:

- maximum gyrotron power, $P_{max} = 3$ MW
- frequency, $f_0 = 110$ GHz
- vacuum wavelength, $\lambda_0 \approx 2.7$ mm

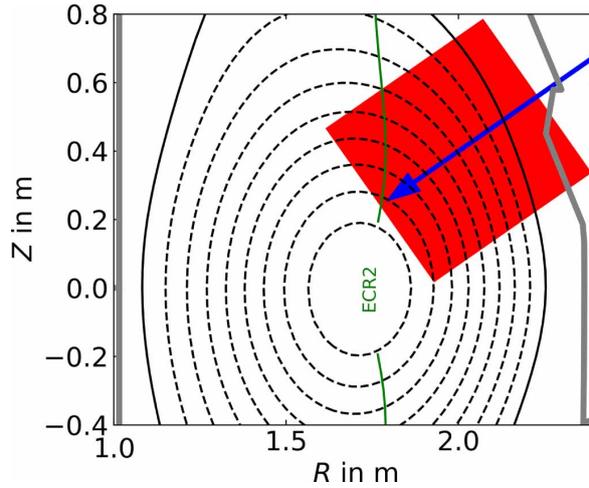


Figure 6.1: Cross-section of the DIII-D tokamak showing microwave launch. The dashed black lines represent closed flux surfaces, the solid black line represents the separatrix, and the thick grey line indicates the vacuum vessel. The injected microwave beam is illustrated by the blue arrow with the 2nd harmonic electron cyclotron resonance indicated by the green line. The red rectangle shows the spatial domain of the 2D full-wave simulations that were carried out to compare to experiment.

- beam waist radius at launcher, $w_0 = 65 \text{ mm} \approx 24\lambda_0$

As the waist of the Gaussian beam is significantly larger than the wavelength, it can be approximated as a plane wave solution. In actuality, the divergence of the beam is $\sim 1^\circ$

In order to find how much the microwave beam is additionally broadened compared to what would be expected from simple refraction due to the plasma density profile, the deposition profile of the beam is first measured from the heating profile. However, this is made more complex by transport in the plasma occurring over the time required to measure the heating profile. This transport results in the width of the measured heating profile appearing larger than the width of the region illuminated by the beam.

In order to isolate the microwave beam deposition profile from the heating profile made to appear larger by transport, the injected microwave beam is power-modulated with a square wave at 70 Hz in an otherwise stationary discharge. This generates a periodic heat wave propagating away from the highly localised deposition layer. The modulation is sufficiently fast that the resulting density perturbation is negligible [81, 82].

The electron temperature is then measured using fundamental ECE with a temporal resolution of $2 \mu\text{s}$ and a spatial resolution of 5 mm. The time variation of the temperature profile then allows the deduction of heat transport coefficients which can be used to calculate the initial deposition profile from the measured heating profile.

Scenario	Shot	I_p	κ	δ	ρ_{ECH}	$\bar{n}_e _{\rho=0.95}$	$\frac{\delta n_e}{n_e} _{\rho=0.95}$	b
diverted L-mode	165 078	1.0	1.80	0.45	0.30	2.9	0.12	2.7
H-mode	165 146	1.0	1.83	0.32	0.25	3.2	0.05	1.7
$-\delta$ L-mode	166 192	0.9	1.32	-0.38	0.45	3.0	0.03	1.5

Table 6.1: Configuration of relevant discharges on DIII-D, all with $B_t = 2$ T. The plasma current I_p is in units of MA and the background plasma density at a normalised radius of $\rho = 0.95$ is in units of $\times 10^{19} \text{ m}^{-3}$. Parameters κ and δ describe the shaping of the plasma, and ρ_{ECH} is the normalised radius of the power deposition location. The broadening factor, b , comes from the width of the deposition profile (after transport analysis) divided by the width predicted by TORAY-GA. This broadening can be seen to scale with fluctuation level $\delta n_e/n_e$ at $\rho = 0.95$.

The ray tracing code TORAY-GA [83] is then used to predict what the power deposition profile should be without additional broadening. This code accounts for the refractive effects of the background plasma, allowing for 1D temperature and density profiles, but not for the effect of plasma density fluctuations. Comparing the measured power deposition profile with these predictions shows that the power deposition is substantially broadened compared to what would be expected from TORAY-GA. This is in agreement with other experiments [84, 85].

By comparing the predicted deposition profile from TORAY-GA with the one measured, a beam broadening factor is deduced. The results for three distinct operating scenarios are given in Table 6.1.

To verify whether the primary cause of this broadening is due to edge turbulence, these experimental results are compared to simulations.

6.3 Simulations

At the plasma edge, fluctuation levels of turbulent density can exceed 100% of the background density [12]. While the fluctuations occur on similar length-scales to microwave wavelength, turbulence in magnetised plasmas is elongated along magnetic field lines. This means that the correlation length of the turbulence in the toroidal direction is very long compared to microwave wavelength, so does not cause significant scattering. Perpendicular to \mathbf{B}_0 the correlation length of the turbulent density structures, L_\perp , typically scales with the drift parameter $\rho_s = \sqrt{T_e m_i}/(eB_0)$ such that $L_\perp \approx 5 - 10\rho_s$ [86]. The resulting turbulent density structure size for a typical L-mode discharge is then $L_\perp \approx 0.5 \text{ cm}$ [86] which is the same

order of magnitude as the vacuum wavelength of the injected microwave beam, resulting in strong scattering [53], though this can increase to of order a few centimetres. The turbulent structure size is typically smaller after the L-H transition. The large fluctuation levels and structure size being comparable to λ_0 means that a full-wave code is required to accurately model the scattering of the microwave beam by turbulence.

To this end, the full-wave cold-plasma code EMIT-2D (as described in Sec. 4.2) was used. As it is a cold-plasma code it does not simulate the absorption of the X-mode at the resonance, merely simulating its propagation to the absorption layer. A 2D version of the code is sufficient as the beam propagates mostly perpendicular to the background magnetic field, similar to the scenario described in [29]. As such, computationally expensive 3D simulations can be avoided.

EMIT-2D allows for the simulation of the propagation of EM waves with arbitrary density and magnetic field profiles. Compared to the speed of microwave propagation, the plasma density fluctuations appear to be ‘frozen’ due to their comparatively low velocity of $\sim 10^4$ ms⁻¹. Therefore, to account for the effect of turbulence on the microwave beam, the microwave is simulated propagating through a series of uncorrelated snapshots, reading out the E_{RMS} at the ‘absorption layer’ for each snapshot, then averaging the read out E_{RMS} over the ensemble to find the ensemble-average E_{RMS} . The broadening is then calculated by fitting a Gaussian to this ensemble-average E_{RMS} and comparing to the Gaussian beam from a background case where no turbulence was present, as shown in Fig. 6.2.

In this example case shown in Fig. 6.2, the width of the Gaussian fitted to the ensemble average can be compared with the background case to give a relative broadening of $w_{\text{fluct}}/w_{\text{background}} \approx 2.1$. Both beams have been normalised to the amplitude of the background beam. It can be seen that the broadening of the beam by turbulence has also decreased the central peak amplitude of the wave.

While this describes the general approach to calculating predictions of beam broadening by plasma density turbulence from simulations, before the exact simulation set-up was decided upon, a number of different effects were considered.

6.3.1 Fluid vs Synthetic Turbulence

As stated above, before I started my PhD, work on this project had been carried out by a previous PhD student using EMIT-3D. The computational cost of a 3D code had forced

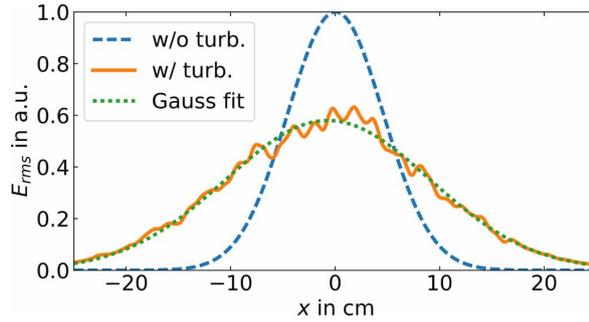


Figure 6.2: E_{RMS} profile of microwave beam at absorption region perpendicular to B_0 and direction of beam propagation. Broadening is calculated by comparing to a background case where no turbulence was present.

them to use a much smaller simulation domain, and they used fluid turbulence generated using the Hermes fluid code in the BOUT++ framework [87]. Given the larger simulation domains I wanted to use, synthetic turbulence offered a computationally cheaper and easier method of turbulence generation. In order to verify that this would not affect our results, I performed a comparison to the fluid turbulence over the initial, smaller simulation domain. Note that this also necessitated using a smaller microwave beam waist of $6\lambda_0$ than was used in the experiments.

The turbulence parameters used were chosen to match the fluid turbulence profiles produced by a previous PhD student. The case for the diverted L-mode was selected for comparison. The correlation lengths were therefore $L_r = 15$ mm, $L_p = 4$ mm. The full background density profile and fluctuation envelope are shown in Fig. 6.3 however a reduced simulation domain was used, only simulating the beam propagating through the turbulent layer at the edge rather than its whole journey to the absorption region. These envelopes were scaled from Hermes outputs to match experimental data. It should be noted that the correlation lengths and fluctuation amplitude envelope used here are not the same as those used in the final comparison to DIII-D. These are used here only because they matched the already existing fluid turbulence profiles so saved in computational expense generating more. This was carried out for a small ensemble of twenty turbulence profiles of each kind to save computational cost, so resulted in larger uncertainties on the broadening than the final results which used a larger ensemble of profiles.

The density fluctuations for the synthetic turbulence were generated using a truncated

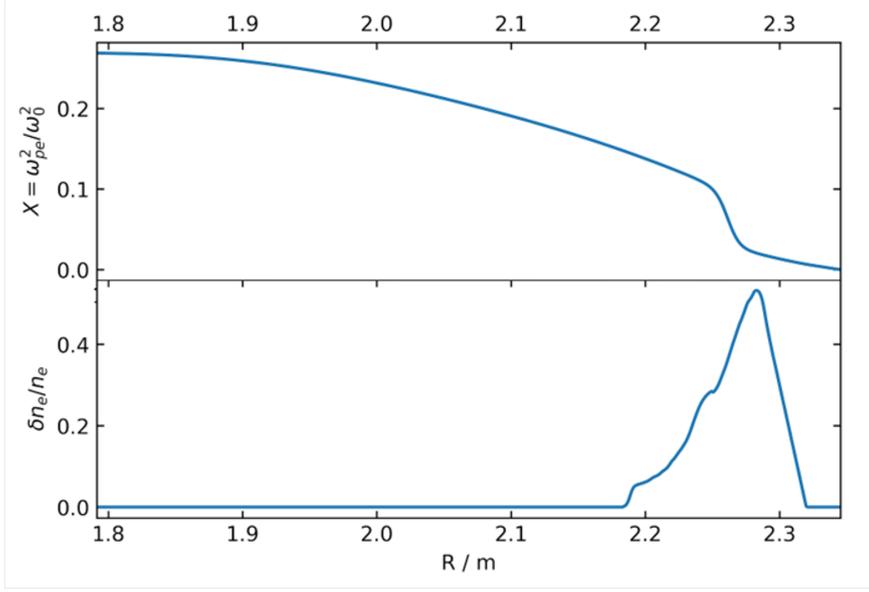


Figure 6.3: Background density and fluctuation envelope for the L-mode case, as was used by for the fluid turbulence vs synthetic turbulence comparison. It should be noted that the actual simulation domain started at the outer edge and is cut off at $R = 2.18$ m.

sum of Fourier-like modes given by [29]

$$\delta n_e(x, z) = \sum_i^{M_i} \sum_j^{M_j} A_{ij} \cos(k_{x,i}x + k_{z,j}z + \phi_{i,j}) \quad (6.1)$$

where A_{ij} are the amplitudes of the modes, and $\phi_{i,j}$ are independent random phases uniformly distributed between 0 and 2π .

The amplitudes are related to the structure size by [29]

$$A_{ij} = \sqrt{\frac{2\sigma_x\sigma_z}{\pi}} \exp[-\sigma_x^2 k_{x,i}^2 - \sigma_z^2 k_{z,j}^2] \quad (6.2)$$

where $\sigma_x = \pi L_x / a_x$ and $\sigma_z = \pi L_z / a_z$, a_x and a_z are the box size, and L_x and L_z are the turbulence correlation lengths in the x and z directions in the simulation domain respectively. In a tokamak scenario, these correspond approximately to the radial and poloidal directions respectively for perpendicular injection along the radial coordinate.

The results for this scenario are shown in Fig. 6.4 and are summarised as follows:

- fluid turbulence gives broadening factor of 1.3 ± 0.1
- synthetic turbulence gives broadening factor of 1.2 ± 0.1

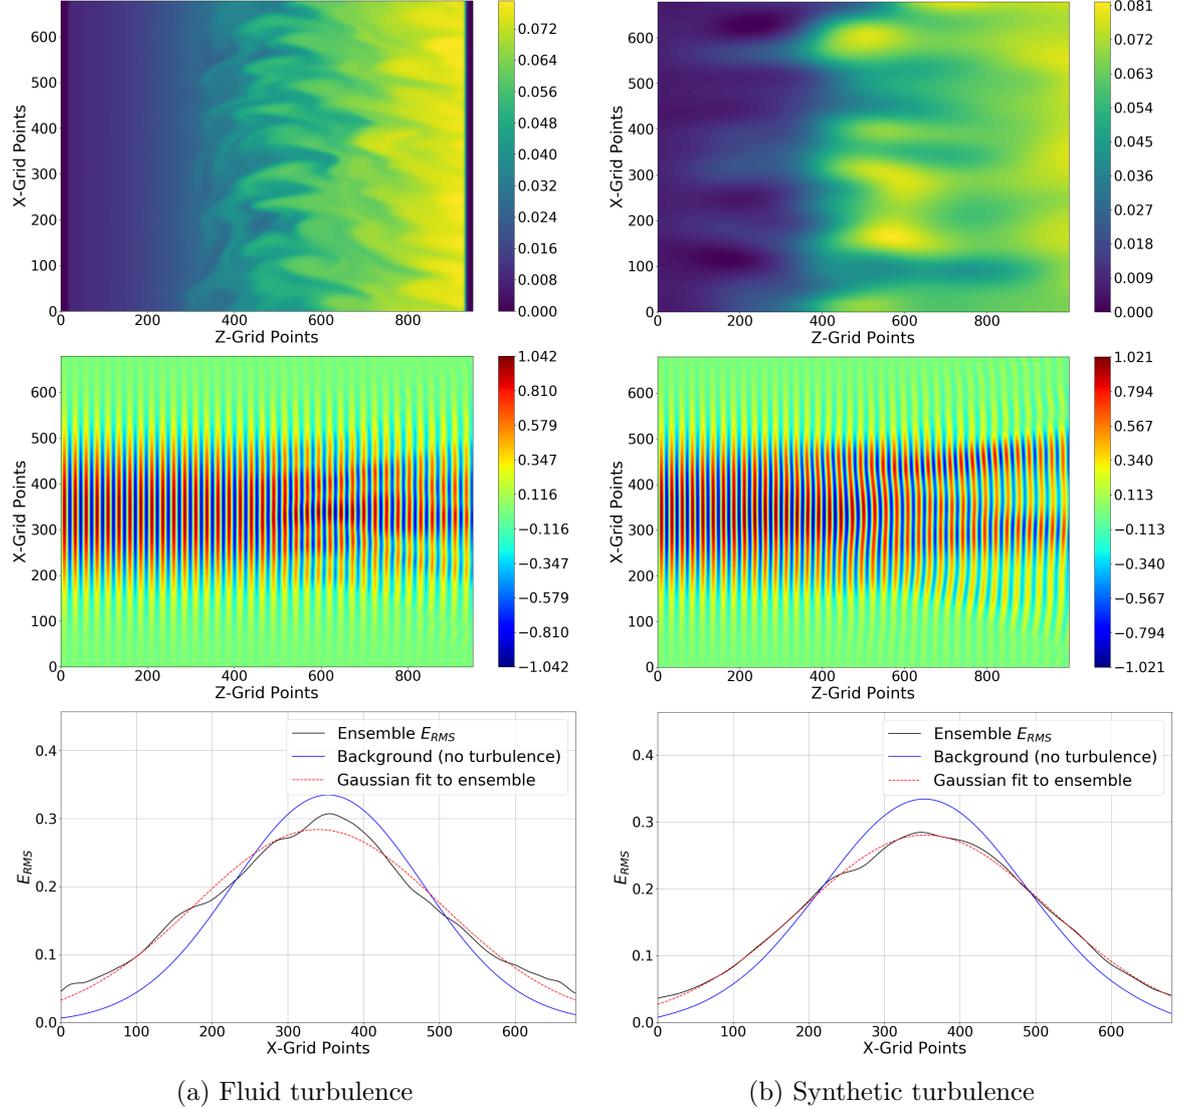


Figure 6.4: Comparing the broadening found running simulations through fluid and synthetic turbulence with the same macroscopic properties of turbulence correlation lengths, background density, and fluctuation level envelope. The top panel shows the normalised density n_e/n_{crit} for a single turbulent snapshot. The middle panel shows a single snapshot of the wave electric field passing through a single turbulent snapshot. The bottom panel shows the RMS electric field read out at the backplane of the simulation. Axis are plotted in grid points where a resolution of 25 grid points per vacuum wavelength was used. For the fluid turbulence, the broadening is found to be 1.3 ± 0.1 and for the synthetic turbulence, it's found to be 1.2 ± 0.1 .

where the uncertainty on the broadening factors is the standard error on the mean across the ensemble average E_{RMS} . Given the good agreement between the two, we decided that synthetic turbulence would be sufficient for further simulations.

This allowed us to generate our own turbulent density profiles with parameters matching experiments which were used for the remainder of this chapter. Using experimental data a correlation length of $L_c = 5$ mm was chosen for both the radial and poloidal correlation length. They were set equal to each other as experimental data only existed for the radial direction. The fluctuation envelope was set by fitting an exponential to the experimental data points before capping the envelope at a fixed value informed by theory [12]. More on the exact simulation setups can be found in Sec. 6.3.3.

6.3.2 Effect of Core Turbulence

Whilst there is a layer of turbulence at the edge of fusion plasmas with high fluctuation levels, turbulence exists throughout the plasma. We anticipated the effect of core turbulence to be much lower due to the small fluctuation level, but were aware that the beam travelled through a greater distance of it all the way to the absorption region, so were concerned that the cumulative effect might be significant. In addition, though the fluctuation level is lower, it is occurring at higher densities so the absolute fluctuation amplitude is potentially significant with the possibility of approaching cut-off densities. To ensure that we were not missing important effects by neglecting the effect of core turbulence with its low fluctuation level, we performed a comparison between a case including and excluding core turbulence for the diverted L-mode scenario. As with the study of fluid turbulence compared to synthetic turbulence, this was only carried out for a small ensemble of twenty turbulence profiles, resulting in larger uncertainties on the broadening than the final results which used a larger ensemble of profiles.

The background density profile and fluctuation envelopes used can be seen in Fig. 6.5. This time, the full beam was simulated with a waist radius of $24\lambda_0$ as was the full propagation to the absorption region. As I didn't have experimental data for the turbulence scale lengths in the plasma core, I set the scale length as constant throughout the plasma at 5 mm.

The results for this scenario are as follows:

- With core turbulence gives broadening factor of 2.3 ± 0.1

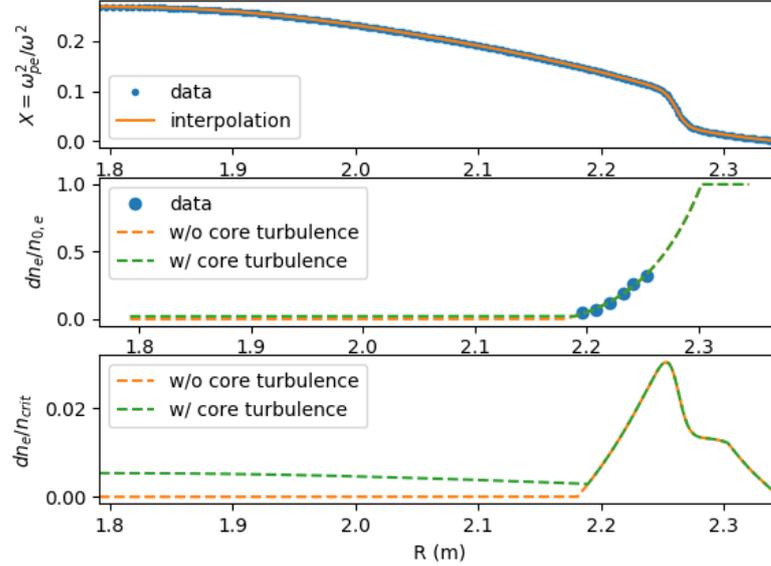


Figure 6.5: Simulation set-up for inclusion of core turbulence to the diverted L-mode case. Top panel shows the background density, middle panel shows the normalised fluctuation level envelope with and without the inclusion of core turbulence at a level of 1% of the background density. The bottom data shows the fluctuation amplitude normalised to the critical density . The experimental data for background density and fluctuation level are shown as dots, with the fits used as inputs for the simulation shown as dashed or solid lines.

- Without core turbulence gives broadening factor of 2.2 ± 0.1

Based on this, from this point on we will neglect the effect of core turbulence. This is in agreement with previous studies [74].

6.3.3 Final Simulation Set-Up

Using the results of the previous sections, we decided to use synthetic turbulence with no turbulence present in the core. The correlation length was set at $L_c = 5$ mm, informed by Doppler reflectometry diagnostic data. The background density profiles were found by performing an interpolation to the experimental data obtained from Thomson Scattering diagnostic measurements. The fluctuation envelope for the turbulence was found by performing an exponential fit to the experimental data obtained from the Beam Emission Spectroscopy measurements, capped at a maximum value informed by the literature [12]. For the diverted L-mode case, this maximum was set at 100%, for the H-mode it was set at 40%, and for the negative triangularity L-mode it was set at 20% of the background density.

Scenario	Resolution	N_x	N_z	CFL	N_t	ν/ω_0
Diverted L-mode	25	5000	5076	0.5	15000	0
H-mode	25	5000	5308	0.5	20000	0
$-\delta$ L-mode	25	5000	4447	0.5	15000	0

Table 6.2: A summary of the numerical parameters used in the benchmark scenarios. Resolution is the number of Yee cells per vacuum wavelength, and N_x and N_z are the number of Yee cells over which the simulation domain extends. CFL relates the time step to the grid spacing, and N_t is the total number of timesteps over which the simulation was run. Finally, ν is the numerical collision frequency outside of the damping boundaries normalised to the vacuum wave frequency.

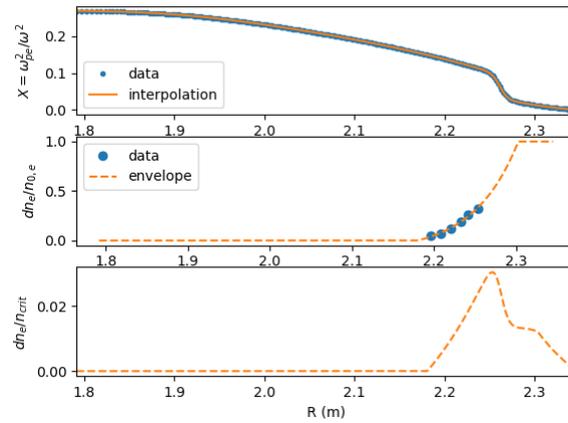
The resultant background density and fluctuation level profiles as used in the simulations are shown in Fig. 6.6. These were used to generate the turbulent density profiles using a synthetic turbulence generator [29] described in more detail in Sec. 7.2.1.

The numerical parameters for the different scenarios are given in Table 6.2. The differing number of grid points in the z -direction are due to the fact that the distance to the backplane of the simulation was set via the distance to the absorption plane from the antenna for the relevant scenario on DIII-D. The number of timesteps was set by waiting at least five full wave-periods after steady-state was reached.

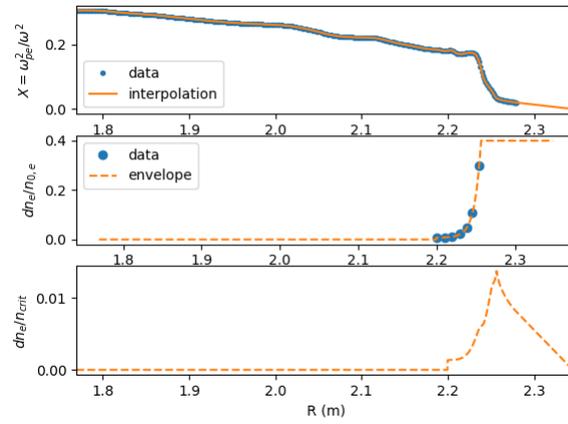
6.4 Results

Before the simulation results could be compared to the experiment, it was important to consider the effect of uncertainty in the diagnostic data. The BES data used to get the fluctuation amplitude of the turbulence had an uncertainty of 4% and a spatial uncertainty of 1 cm. The Thomson Scattering data used to get the background density profiles also had a spatial uncertainty of 1 cm.

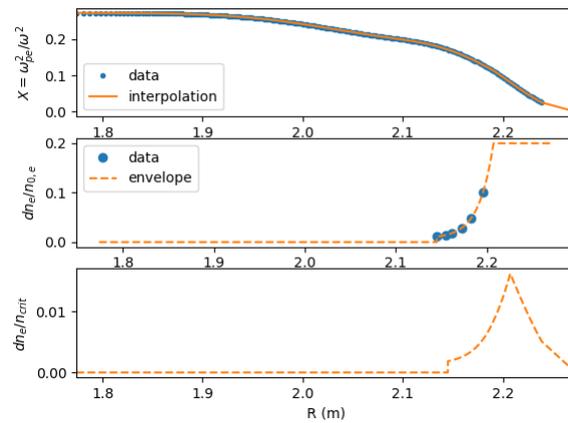
In order to determine the sensitivity of the results to this uncertainty, extremal cases were considered. The case expected to yield the least broadening would be where fluctuation amplitudes were reduced by 4% and pushed into a region of lower density by shifting their spatial location outward by 1 cm while the background density measurements were shifted inward by 1 cm. The case expected to yield the most broadening would be where fluctuation amplitudes were increased by 4% and pushed into a region of higher density by shifting their spatial location inward by 1 cm while the background density measurements were shifted outward by 1 cm. The results are summarised in Tab. 6.3.



(a) Diverted L-mode



(b) H-mode



(c) Negative Triangularity L-mode

Figure 6.6: Background density and fluctuation level envelopes for all three scenarios, as used to compare to experiments. The experimental data used to create these profiles are shown as dots.

Scenario	Exp	Min Sim	Default Sim	Max Sim	Fin Sim
div L-mode	2.7 ± 0.3	1.91 ± 0.04	2.27 ± 0.05	2.71 ± 0.07	2.3 ± 0.5
H-mode	1.7 ± 0.2	1.115 ± 0.009	1.53 ± 0.02	2.03 ± 0.04	1.6 ± 0.5
$-\delta$ L-mode	1.4 ± 0.2	1.113 ± 0.009	1.21 ± 0.01	1.30 ± 0.02	1.2 ± 0.2

Table 6.3: Comparison of experimentally measured (Exp) and simulated broadening. Uncertainty on simulated results comes from the standard error on the mean from the ensemble of turbulence profiles for each scenario. The ‘Default Sim’ is the simulated broadening when using the density and fluctuation data as measured rather than at the extremal ends of the possible range. Uncertainty in the diagnostic data used as input for the simulations is used to determine the minimum (Min Sim) and maximum (Max Sim) possible simulated broadening. This leads to large uncertainty on the final simulated (Fin Sim) result, calculated as the midpoint of the range of results, but ultimately yields agreement with the experimental results.

Good agreement between simulation and experiment is seen. However, even very small uncertainties in the diagnostic data used as input for the simulations can create large uncertainties in the simulated broadening. This is particularly the case for the H-mode scenario, where the steep density gradient means that even a small change in the spatial location of the measurements for background density and fluctuation level has a large impact on the simulated broadening. This motivates the need for ever-improving turbulence diagnostics with even better resolution if we want to truly be able to predict this effect accurately.

6.5 Conclusions and Further Work

We have shown that edge turbulence can dramatically broaden the deposition profile of the microwave beam in DIII-D. Experimentally measured beam broadening was compared to first-principle full-wave simulations for a range of scenarios for the first time. As expected, scenarios with larger turbulence fluctuation amplitudes led to more broadening, however, as numerous parameters changed across the scenarios, no correlation can be deduced. The affect of varying fluctuation amplitude is explored more fully in the next chapter. Quantitative agreement between simulations (where the only possible cause of the broadening was scattering by turbulence) and the experimental measurements has proven that turbulent scattering is likely the primary cause for the beam broadening observed on DIII-D. It has also validated a quantitatively accurate predictive simulation tool that can be used to predict beam broadening on future devices. However, the work has also highlighted the significant uncertainty in simulated broadening resulting from uncertainty in diagnostic data used as

input for the simulations. This motivates the need for ever-improving turbulence diagnostics to make our simulations as accurate as possible. It also motivates the need for accurate turbulence modelling when considering predictions for new machines. The effects of variations in other turbulence and plasma parameters such as turbulence scale lengths and background density are explored in the next chapter.

Chapter 7

Parametric Dependence of Beam Broadening by Turbulence

This work was the subject of a paper published at L.A. Holland *et al* 2023 *Nucl. Fusion* **63** 056013 DOI: 10.1088/1741-4326/acc25e [88]. As such, much of the following chapter is reproduced directly with the permission of Nuclear Fusion and EUROfusion.

7.1 Introduction

As stated in previous chapters, microwaves are often used to inject power into MCF plasmas. In magnetically confined fusion (MCF) devices, microwaves are often used to inject power into the plasmas for the purposes of global heating [19, 21], non-inductive start-up [23, 24], and current drive [20]. Localized electron cyclotron current drive (ECCD) can also be used to stabilize MHD instabilities such as neoclassical tearing modes (NTM), requiring precise targeting of microwave beams [89].

In order to inject power, microwave beams must cross the plasma boundary, a region where density fluctuation levels can reach 100 % of the background density on length scales comparable to that of the microwave wavelength [12]. These fluctuations cause scattering of incident microwaves, leading to the broadening of microwave beams travelling through the plasma. For ITER-like scenarios, it has been predicted that this broadening could result in the doubling of the beam width compared to if no turbulence was present [70]. In the previous chapter, we found significant beam broadening on DIII-D, where the likely primary cause was turbulent scattering. For L-mode, the beam width at the deposition region was

more than doubled by the presence of turbulence, and for H-mode, where NTM stabilisation using targeted, localised ECCD will likely be of great importance, density fluctuation were found to result in a beam broadening factor of roughly 1.6. This broadening has the potential to significantly impact the efficiency of both global and local power injection via microwave beams. It's therefore important to be able to predict this effect when designing and running microwave power injection systems. Whilst full-wave simulations using codes like EMIT-2D can achieve this, their computational expense makes them impractical when a wide range of plasma and launch scenarios need to be considered, such as in an optimisation study of a new heating and current drive system. For example, simulating a beam propagating from the antenna to the absorption region at a resolution sufficient to capture the density fluctuations takes around 10 core hours (running for 1 hour over 10 cores). Repeating this for an ensemble of 20 turbulent density profiles would therefore take 200 core hours. Therefore, each new scenario with a different background density profile, different injection angle, different beam width, different microwave frequency, would each take 200 core hours to predict the broadening for. It is therefore important to find ways to predict this effect that allow it to be accounted for in integrated modelling and optimisation studies for future tokamaks in a way that is computationally cheap.

One such approach is the development of analytical models for both O-mode [28] and X-mode [90]. This approach has the benefit of allowing for very fast prediction of beam broadening based on plasma and beam parameters, making it ideal for optimisation studies or inter-shot analysis. However, these models utilise eikonal methods based on assumptions that the fluctuation level is small, which will not be applicable in the case of strong turbulence locally creating plasma densities close to or even above the cut-off density in the wave's path, meaning that they will not always be applicable for some fusion-relevant scenarios.

Ray tracing methods where the scattering effect is described using a Fokker-Planck solver [27, 70, 91] are another alternative. They are less quick than analytical methods as they still require simulation, but are significantly less computationally expensive than full-wave codes. There have also been studies using a beam tracing code based on the wave kinetic equation, incorporating the effect of turbulence via a scattering operator derived under the Born approximation [92]. However, these methods are also valid only within certain limits, such as when the turbulence amplitude is small compared to the cut-off density, or when the refractive index does not vary significantly over length scales that are small compared to the

wavelength.

In order to simulate particular fusion-relevant scenarios where the turbulent scale length is comparable to the microwave wavelength and the fluctuation level is high, directly simulating the beam path through turbulent plasma using full-wave codes is often required. Using this method, 1D parameter scans have previously been carried out [53]. However, the range of parameters is yet to be extended to encompass certain fusion-relevant scenarios. Furthermore, it has not been previously investigated if the impact of changing one parameter is independent of changes in other parameters i.e. if the dependencies are ‘separable’ as defined in Sec. 7.2.5.

The work presented here investigates how the broadening of microwave beams by a layer of turbulent plasma depends on plasma and beam parameters. The parameter ranges were set to cover fusion-relevant scenarios such as beams used for ECCD, including ranges where turbulence scale length is comparable to microwave wavelength, and fluctuation level is as high as 50%. Fluctuation level wasn’t raised to 100% as, though this value is reached in fusion plasmas, it is usually only for a narrow width at the plasma edge in low density, whereas these parameter scans will have the turbulence present for a range of thicknesses at a range of background densities which would not be applicable and would result in saturated broadening where the wave signal at the back-plane of the simulation is nearly indistinguishable from background noise. We considered pairwise combinations of parameters to determine whether the dependence on each parameter is independent of the others and conducted a point-wise fit to the data set in order to determine an empirical formula for the beam broadening. By determining the dependence of broadening on turbulence and beam parameters in fusion-relevant scenarios, a predictive model can be developed which does not require full-wave simulations but is still applicable in parameter regimes that are not analytically tractable. This would be of great use in the development of future tokamaks by allowing the quick prediction of beam broadening and could be incorporated into an integrated model used to optimise heating and current drive efficiency. It could also be used to make predictions on timescales useful for inter-shot analysis as well as real-time predictions based on plasma measurements.

7.2 Simulation Set Up

The code EMIT-2D as described in Chapter 4 was used to carry out the simulations for the parameter scans.

By setting \mathbf{B}_0 to point in the y -direction, we launched an X-mode beam in vacuum. We kept the normalized field strength constant throughout the simulation domain at $Y = \omega_{ce}/\omega = 0.5$. Whilst this may not be fully realistic for a fusion reactor, we wanted to limit the number of variables that we would change for this study for the sake of feasibility, and we anticipate that the field strength would not change significantly over the region where turbulence is strongest in the plasma edge.

7.2.1 Turbulent Density Profiles

We chose to study the influence of the following key parameters: background density, fluctuation level, turbulence structure size, width of the turbulence layer, and beam waist. The fluctuations were only in the density profiles, making the turbulence purely electrostatic.

To launch the beam in vacuum, we transitioned the density from zero to its constant background value via a hyperbolic tangent function. In the constant background density region, we transitioned the fluctuation level from zero to a constant value via the same hyperbolic tangent function, before transitioning it back to zero via the same hyperbolic tangent function. Though this sharp increase in density is not necessarily realistic for fusion plasmas, we decided that introducing a pedestal-like background density would increase the dimensionality of the problem too much. As such, this simplified approach was taken. By considering the simulations run through background density profiles, looking at the wave electric field across the whole simulation domain, we were able to determine that this sharp increase in density did not perturb the beam as it remained Gaussian in nature as it propagated through the domain, without a significant or sudden change where it entered the plasma. An example of the background density profile and fluctuation envelope, with an image of a turbulent snapshot and beam propagating through that snapshot, are shown in Fig. 7.1.

The simulation domain shown in Fig. 7.1 is in 2D. This is an acceptable simplification as the turbulent structures are elongated along field lines [12], making the scattering in the toroidal direction small compared to the radial and poloidal direction. The simulation

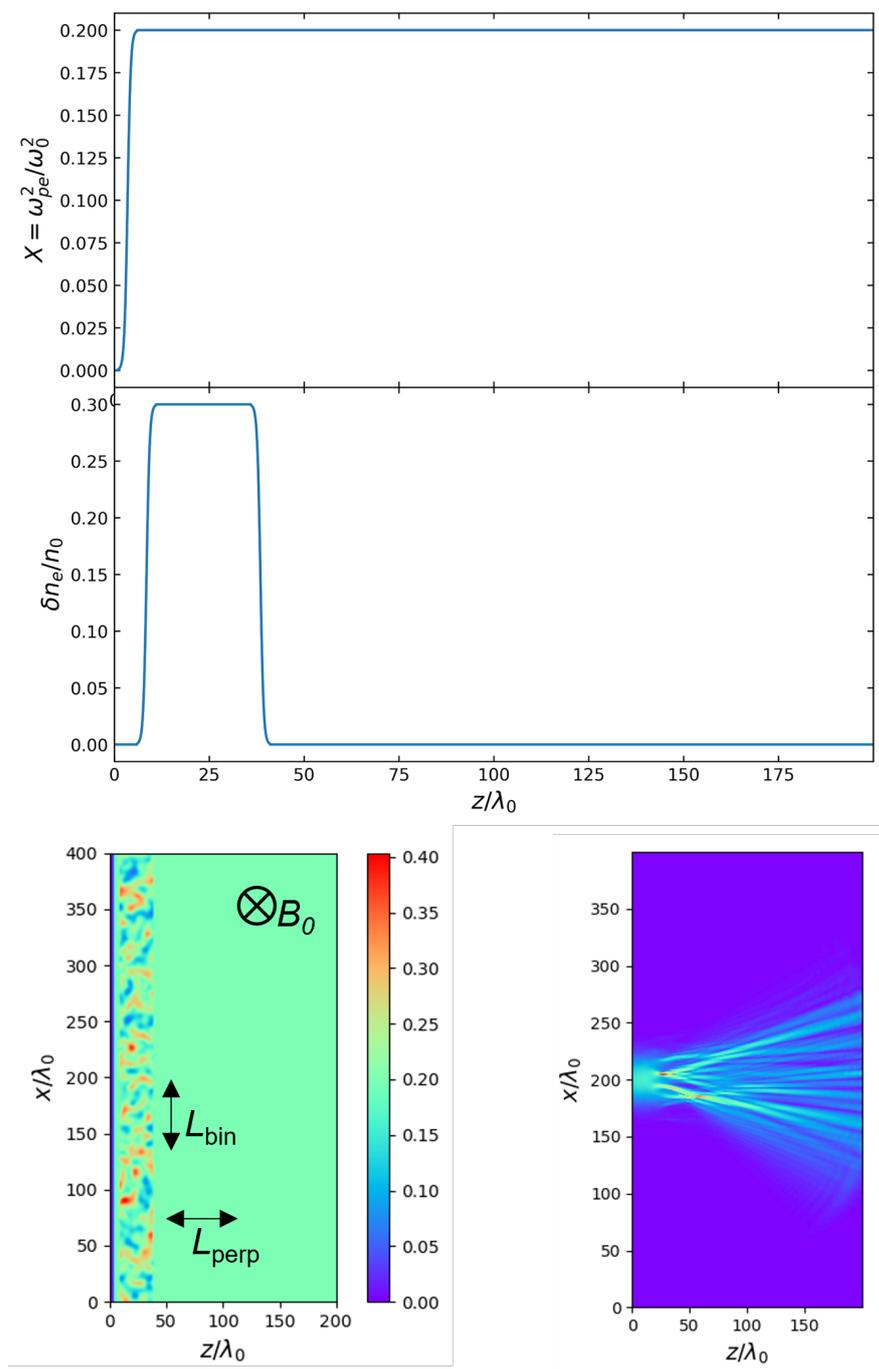


Figure 7.1: Example ensemble average background density profile and fluctuation level envelope (top), an example turbulent density profile (bottom left) and the RMS Electric field (bottom right) of a microwave beam propagating through it. The simulation domain is given in vacuum wavelengths. The colour represents the normalised density $X = \omega_{pe}^2/\omega^2 = n_{e,0}/n_{\text{crit}}$ (bottom left) and the RMS field strength in arbitrary units (bottom right). Arrows are marked on to indicate the direction of the background field as well as the direction in which the two different turbulence correlation lengths are defined.

domain is therefore approximately equivalent to a poloidal cross-section in a tokamak.

The turbulent density profile shown in Fig. 7.1 represents a snapshot of the plasma. In reality, the plasma would be moving at a speed significantly slower than the wave speed. This allows us to make the assumption that the turbulence is ‘frozen’, and the overall effect on the beam can be calculated by averaging over an ensemble of turbulence profiles.

In order to generate an ensemble of turbulent density profiles, we used synthetic turbulence. This allowed full control over the turbulence parameters and, for the large number of profiles needed, was less computationally demanding than fluid turbulence generation from a turbulence code.

The turbulent density profiles were generated using the same method described in Sec. 6.3.1, repeated here for convenience. The density fluctuations were generated using a truncated sum of Fourier-like modes given by [29]

$$\delta n_e(x, z) = \sum_i^{M_i} \sum_j^{M_j} A_{ij} \cos(k_{x,i}x + k_{z,j}z + \phi_{i,j}) \quad (7.1)$$

where A_{ij} are the amplitudes of the modes, and $\phi_{i,j}$ are independent random phases uniformly distributed between 0 and 2π and M_i and M_j are the number of harmonics in x and z respectively.

The amplitudes are related to the structure size by [29]

$$A_{ij} = \sqrt{\frac{2\sigma_x\sigma_z}{\pi}} \exp[-\sigma_x^2 k_{x,i}^2 - \sigma_z^2 k_{z,j}^2] \quad (7.2)$$

where $\sigma_x = \pi L_{\text{bin}}/a_x$ and $\sigma_z = \pi L_{\text{perp}}/a_z$, a_x and a_z are the box size, and L_{bin} and L_{perp} are the turbulence correlation lengths in the x and z directions respectively. The number of harmonics (M_i and M_j) in Eq. (7.1) are set such that $\exp[-\pi^2 M_{x,z}^2 (L_{\text{bin,perp}}/a_{x,z})^2] < 1 \times 10^{-8}$. Within the geometry of the simulations, the correlation length perpendicular to the beam path is L_{perp} , and the correlation length bi-normal to the beam path and the background magnetic field is L_{bin} . In a tokamak scenario, these correspond approximately to the radial and poloidal correlation lengths respectively for perpendicular injection along the radial coordinate.

We generated an ensemble of twenty turbulence profiles for each combination of turbulence correlation lengths. We then scaled the turbulence to the required fluctuation am-

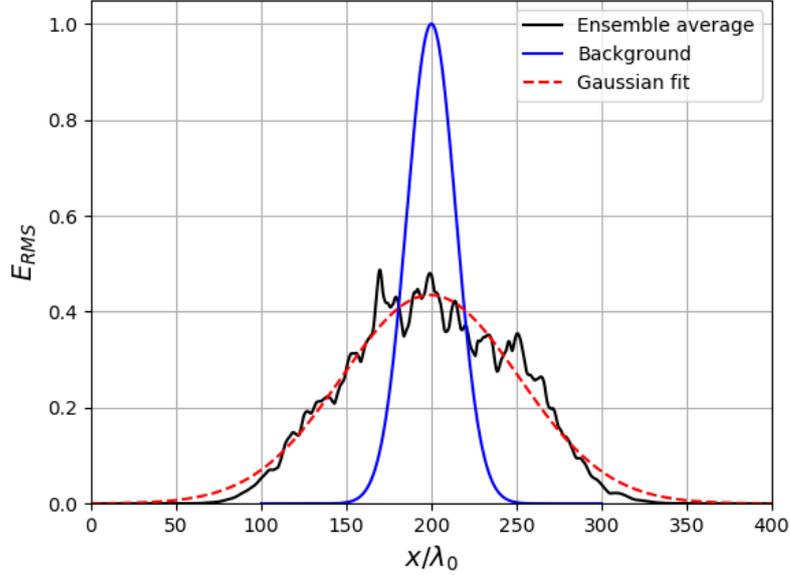


Figure 7.2: An example of the RMS electric field at the back-plane of the simulation domain, after travelling through the turbulent layer, averaged over an ensemble of 20 uncorrelated turbulence profiles. This is the "base case" as defined in Table 7.1. A Gaussian fit to the ensemble average is performed and compared to the background Gaussian beam to find the broadening factor.

plitude, applied an envelope function to achieve the required turbulence layer thickness, and added it to the background density. Any negative density values were truncated to zero, resulting in the negative fluctuation amplitude being slightly lower than the positive fluctuation amplitude.

7.2.2 Calculating Beam Broadening

In order to determine the broadening, Γ , for each combination of parameters in a scan, we found the ensemble average RMS electric field (E_{RMS}) at the back-plane of the simulation for each turbulent snapshot. We then calculated the ensemble average of these E_{RMS} profiles for the ensemble of turbulence snapshots and compared it to a beam which propagated through an equivalent background profile with no turbulence present. For all simulations, we used an ensemble size of 20 turbulence profiles. This was sufficient for the centre of the broadened beam to align closely with the background case. An example of this is shown in Fig. 7.2. This approach results in a statistical uncertainty on the broadening value given by the standard error of the ensemble.

Table 7.1: Parameters that are varied, how they are varied, and the ‘base’ value each parameter takes in scans where it is not varied. Note that the background density is only varied through nine values instead of ten, as at the timescales needed for higher densities, the stability condition proved prohibitively computationally expensive.

Parameter	Varied as	Min	Max	Step	Base
Perpendicular correlation length	$\log(L_{\text{perp}}/\lambda_0)$	-1.5	3.0	0.5	2.0
Binormal correlation length	$\log(L_{\text{bin}}/\lambda_0)$	-2.5	2.0	0.5	2.0
Beam waist radius	w_0/λ_0	2.5	25.0	2.5	20.0
Fluctuation level	$\delta n_e/n_{e,0}$	0.05	0.50	0.05	0.30
Background density	$X = n_{e,0}/n_{\text{crit}}$	0.05	0.45	0.05	0.20
Width of turbulence layer	$W_{\text{turb}}/\lambda_0$	5	50	5	30

7.2.3 Parameter Scans

The parameters we investigated are listed in Table 7.1, with the value they take in the ‘base case’ included in every parameter scan as well as the ranges they are varied over. All parameters are normalised. Length scales are normalised to the vacuum wavelength, λ_0 , background density is normalised to the O-mode cut-off density n_{crit} , and fluctuation amplitude is normalised to background density. We varied the turbulence correlation lengths logarithmically, as we anticipated that broadening will be much smaller when the ratio between correlation length and wavelength becomes either small or large, as seen in previous work [53]. We, therefore, wanted to treat the ratio symmetrically, which is done by varying its logarithm.

From previous work in the field, we expected scattering to be maximal when the correlation lengths are close to the vacuum wavelength, and to decrease as the correlation length moves from this value in either direction [53].

For low-density gradients within the perturbed density (meaning that fluctuation level is small or correlation length is large compared to wave length such that any density gradients are low over wavelength length scales), we expected scattering to increase quadratically with fluctuation amplitude [27, 53, 90]. We also expected it to increase with background density and width of the turbulence layer [53, 90]. For the microwave beam waist, we expected an initial increase followed by a gradual decrease, as seen in previous studies [53]. This is due to the dependence of the microwave beam width on its initial waist varying similarly non-monotonically, as can be seen from Eq. 3.31.

In order to determine the separability of the dependence on each of these parameters, we carried out a series of 2D scans considering each of the pairwise combinations. Each of

these 2D scans required 100 ensembles of 20 simulations to be run, each at an approximate computational cost of 80 CPU hours.

7.2.4 Numerical Parameters

For the base case in the parameter scans, the numerical parameters used were as follows:

- Resolution = 25 Yee cells per vacuum wavelength
- $N_x = 10000$ Yee cells in the x -direction
- $N_z = 5000$ Yee cells in the x -direction
- $CFL = 0.5$
- $N_t = 14000$ timesteps = 280 vacuum wave periods
- $\nu/\omega_0 = 0$

These parameters were also used for simulations with the same background density as the base case, unless one of the turbulence correlation lengths was below $1 \lambda_0$ in which case a resolution of 50 Yee cells per vacuum wavelength was used and the number of timesteps doubled accordingly so that the total simulation time remained the same number of vacuum wave time periods.

The only parameter to change with changing background density was the number of timesteps that the simulation was run for. Higher densities require the simulation to run for longer before it reaches steady-state. These are given in Table 7.2 for the case of a resolution of 25 Yee cells per wavelength. When either turbulence correlation length was smaller than one wavelength and a resolution of 50 Yee cells per vacuum wavelength was used, the number of timesteps was doubled accordingly so that it remained the same number of vacuum wave periods.

7.2.5 Fitting

As before, we define broadening as $\Gamma = \text{FWHM}_{\text{ensemble}}/\text{FWHM}_{\text{background}}$ where $\text{FWHM}_{\text{ensemble}}$ is the full width at half maximum of the gaussian fitted to the E_{RMS} at the backplane of the simulation, averaged over the ensemble of simulations through the ensemble of turbulence

Table 7.2: Number of timesteps that the simulation was run for based on background density, assuming a resolution of 25 Yee cells per wavelength. Higher background densities result in the simulation taking longer to reach steady state. In all cases, the code was run for at least one full wave period after steady state was reached.

$n_{e,0}/n_{\text{crit}}$	N_t
0.05	12000
0.10	12500
0.15	13000
0.20	14000
0.25	15000
0.30	17000
0.35	19000
0.40	22500
0.45	28000

profiles, and $\text{FWHM}_{\text{background}}$ is the full width at half maximum of the E_{RMS} at the back-plane of a simulation of an equivalent plasma without the turbulence present. To determine the dependence of Γ on each parameter and confirm the separability of the dependencies, we performed point-wise fits to the data-set from each 2D parameter scan.

By separable dependence, we mean that, for two parameters a and b , the overall broadening can be expressed as a product of two independent functions, f and g , such that

$$\Gamma = 1 + f(a)g(b). \quad (7.3)$$

Within the context of a point-wise fit, the broadening for a particular combination of parameters can then be calculated as

$$\Gamma(a_i, b_j) = 1 + Cf(a_i)g(b_j) = 1 + Cf_i g_j. \quad (7.4)$$

where C is a normalisation factor, fixing the value of the ‘base case’ broadening to be $\Gamma_{\text{basecase}} = 1 + C$. To allow for the possibility that the dependence is not separable, we included a rotation of the coordinate axis in the fit, adding an additional fit parameter, θ . The fit parameters f_i and g_j then depend on $(a \cos \theta - b \sin \theta)$ and $(a \sin \theta + b \cos \theta)$ as

$$\Gamma(a_i, b_j) = 1 + Cf(a_i \cos \theta - b_j \sin \theta)g(a_i \sin \theta + b_j \cos \theta) = 1 + Cf_{i,j}g_{i,j}. \quad (7.5)$$

As each parameter was varied across ten values, for each 2D parameter scan this resulted in 20 fitting parameters for 100 data points. The exception is scans where background den-

sity is varied where, due to computational constraints, only nine values were included. This is owing to the fact that the size of the simulation domain required increases for increased broadening and that the wave speed in higher-density plasma decreases, making the simulations take significantly longer to run (of order 100 core hours for each individual simulation, therefore 2000 core hours for an ensemble of 20 turbulence profiles). It would be possible to expand upon the range included here, but for this work, they were deemed computationally too expensive.

Once a θ value was found for each 2D scan, we knew which dependencies were separable and which were not. A pair of parameters we deemed to be separable if the θ found for them was within error of zero. We then performed a point-wise fit to the whole data set, only allowing for inseparability where it had been found.

7.3 Parametric Dependence

Performing the fits to the 2D scans individually, the only pairwise combinations where we found the dependence to be inseparable were those for $L_{\text{perp}}/\lambda_0$ vs L_{bin}/λ_0 , and for $\delta n_e/n_{e,0}$ vs $n_{e,0}/n_{\text{crit}}$.

An example plot of a 2D dependence of Γ on a pair of separable parameters (w_0/λ_0 vs $W_{\text{turb}}/\lambda_0$) is shown in Fig. 7.3. Plots of the 2D dependence of Γ on both pairs of inseparable parameters ($L_{\text{perp}}/\lambda_0$ vs L_{bin}/λ_0 and $\delta n_e/n_{e,0}$ vs $n_{e,0}/n_{\text{crit}}$) are shown in Figs. 7.5 and 7.4 respectively. The rest of the 2D scans are shown in Figs. (7.6-7.17) for completeness. In the case of the two fully separable parameters, w_0/λ_0 and $W_{\text{turb}}/\lambda_0$, the same trends can be seen across all 2D scans. For the case of the parameters which have one inseparable dependence, the trends differ when shown in a 2D scan with their partner parameter as opposed to in a 2D scan with a parameter from which they are separable, as would be expected.

The results of a global pointwise fit, as described in Sec 7.2.5, are shown in Fig. 7.19. These show how the broadening depends on each parameter (or linear combination of parameters in the cases where the dependence is not separable). It should be noted that for the linear combinations of parameters, the extremal points only correspond to one ensemble of simulations (for example, the point with the lowest L_{bin} and the highest L_{perp}). Every other point corresponds to multiple ensembles. This is why the errors on the extremal points are noticeably larger. In future studies, it would be possible to add more data points into

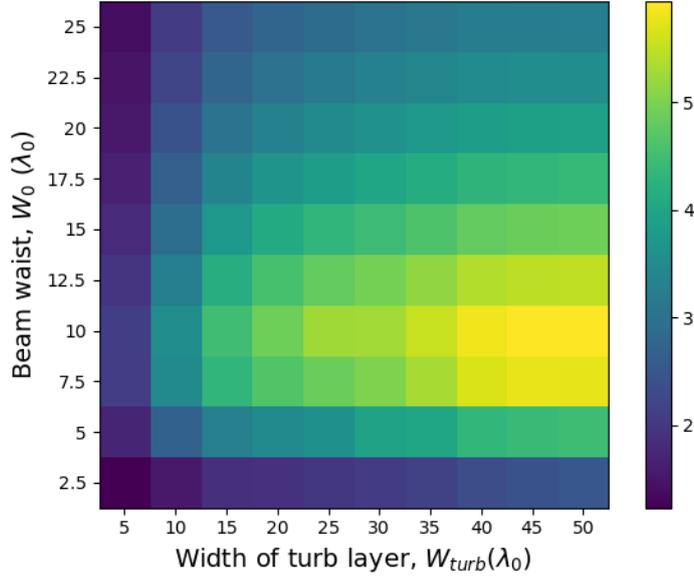


Figure 7.3: Dependence of broadening on microwave beam waist (w_0) and width of the turbulence layer (W_{turb}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

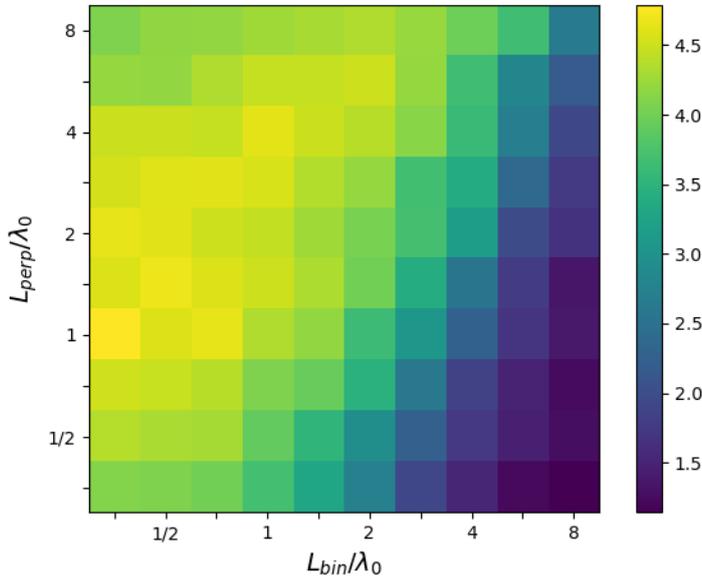


Figure 7.4: Dependence of broadening on turbulence correlation length in the direction perpendicular to the magnetic field (L_{perp}) and in the direction binormal to both the magnetic field and direction of beam propagation (L_{bin}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present. From this, we see that both correlation lengths affect broadening differently, and the tilted nature indicates the dependence is not separable.

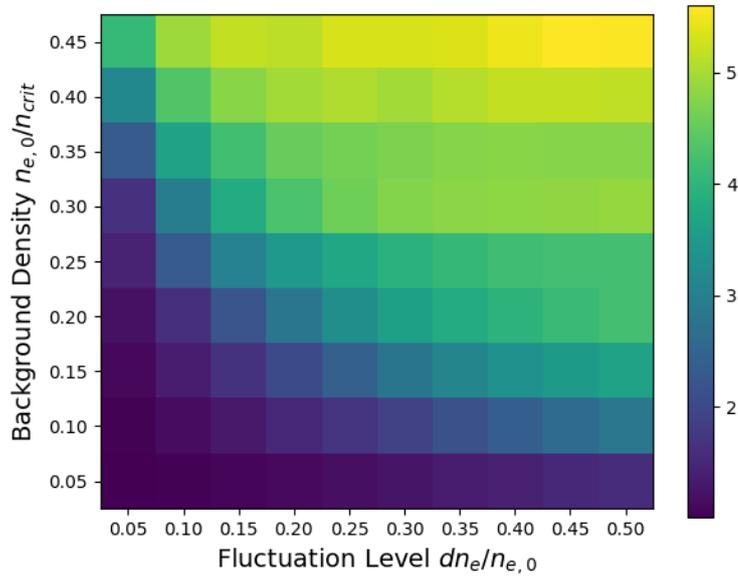


Figure 7.5: Dependence of broadening on background density ($n_{e,0}/n_{crit}$) and fluctuation amplitude ($\delta n_e/n_{e,0}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present. From this, we see that broadening increases with both background density and fluctuation amplitude, and the tilted nature indicates the dependence is not separable.

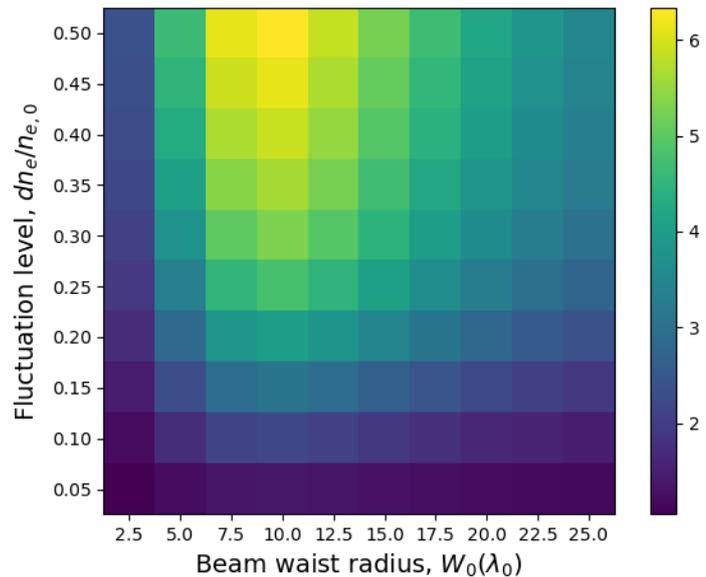


Figure 7.6: Dependence of broadening on microwave beam waist (w_0) and fluctuation level ($dn_e/n_{e,0}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

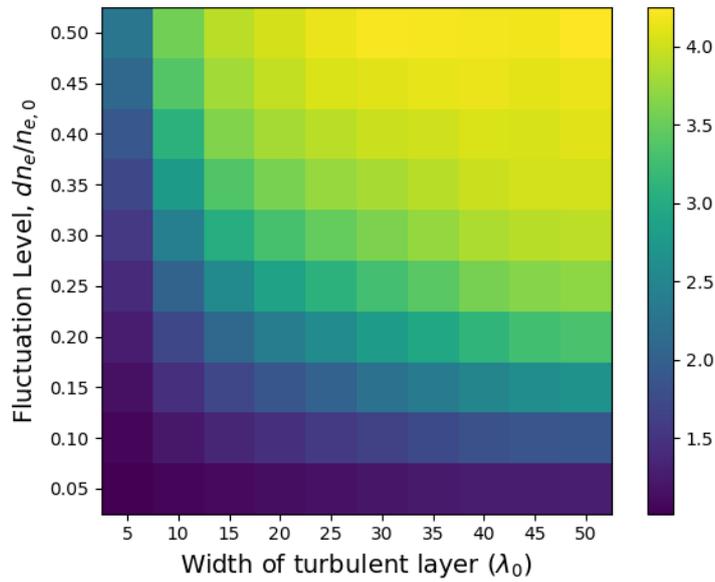


Figure 7.7: Dependence of broadening on width of turbulence layer (W_{turb}) and fluctuation level ($dn_e/n_{e,0}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

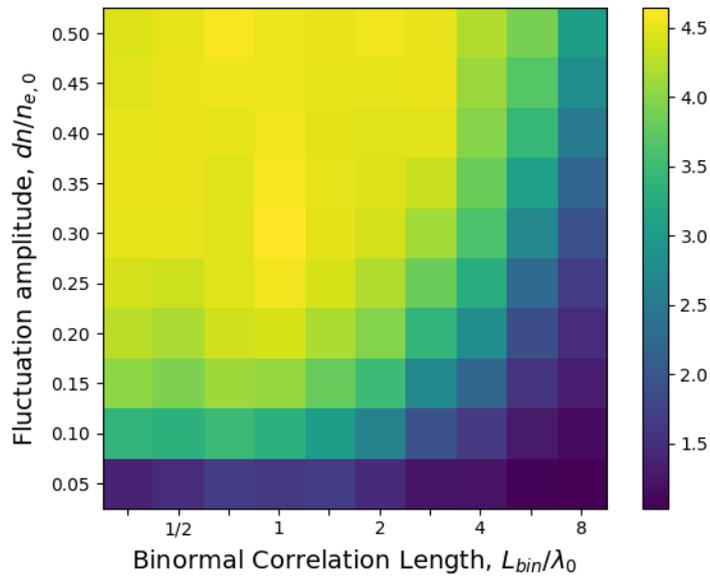


Figure 7.8: Dependence of broadening on binormal correlation length (L_{bin}) and fluctuation level ($dn_e/n_{e,0}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

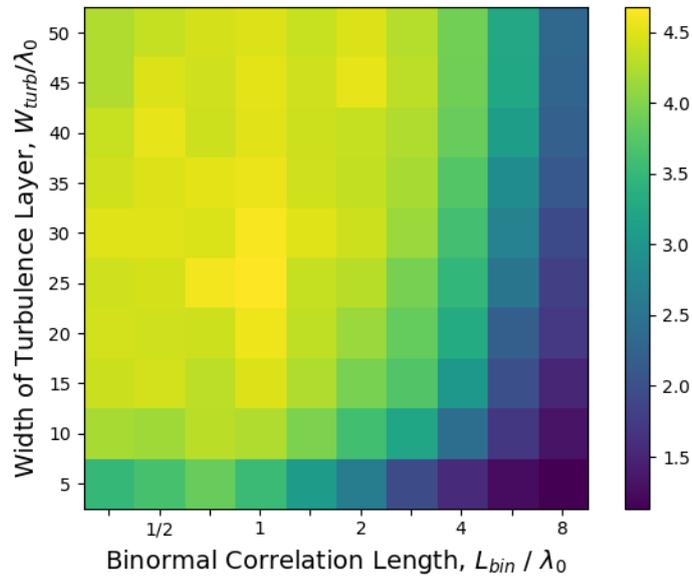


Figure 7.9: Dependence of broadening on binormal correlation length (L_{bin}) and width of turbulence layer (W_{turb}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

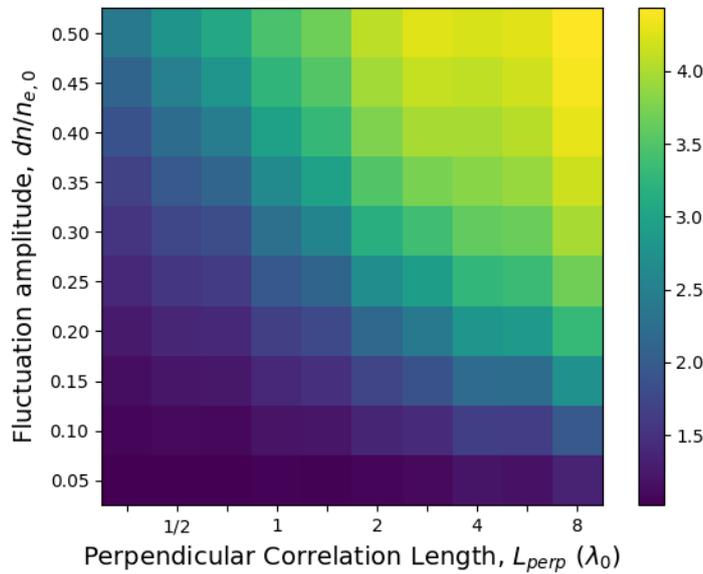


Figure 7.10: Dependence of broadening on perpendicular correlation length (L_{perp}) and fluctuation level ($dn/n_{e,0}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

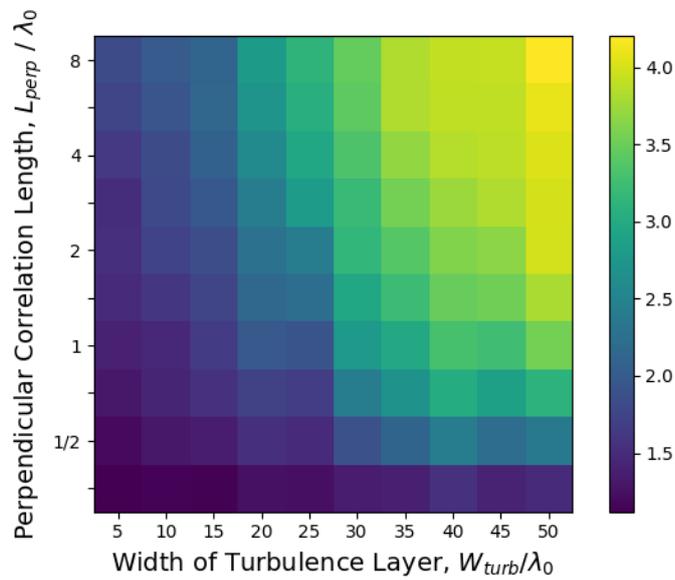


Figure 7.11: Dependence of broadening on perpendicular correlation length (L_{perp}) and width of turbulence layer (W_{turb}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

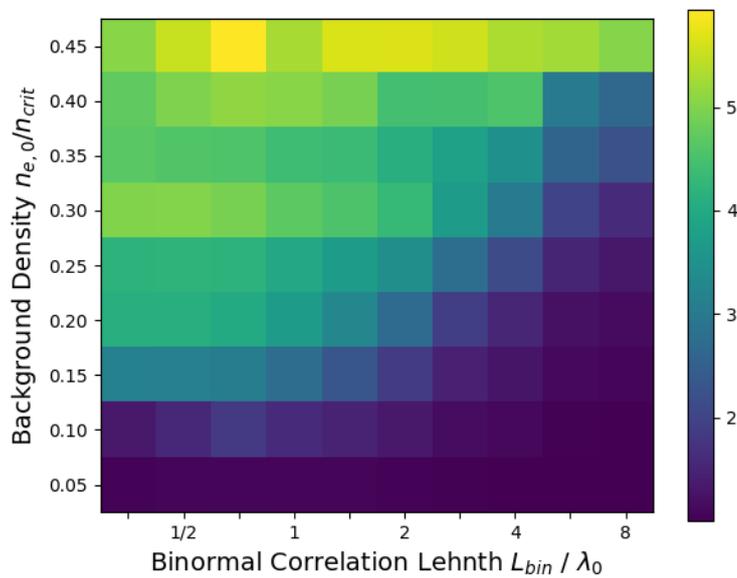


Figure 7.12: Dependence of broadening on binormal correlation length (L_{bin}) and background density ($n_{e,0} / n_{\text{crit}}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

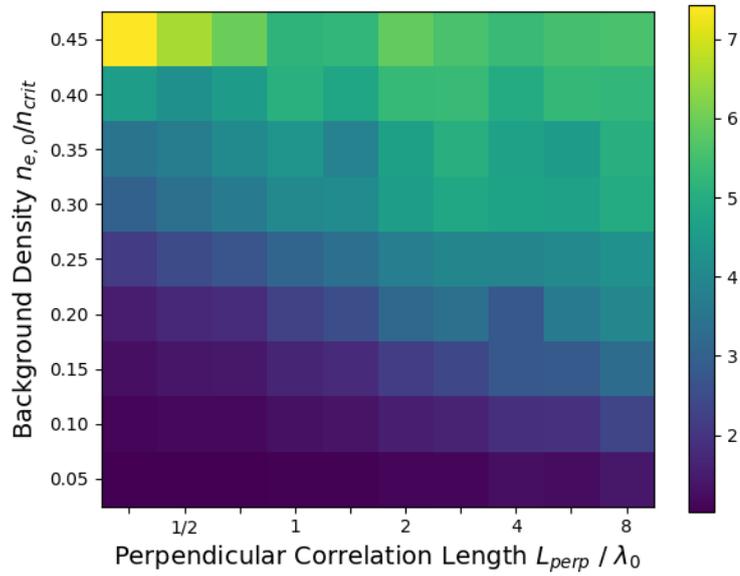


Figure 7.13: Dependence of broadening on perpendicular correlation length (L_{perp}) and background density ($n_{e,0}/n_{\text{crit}}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

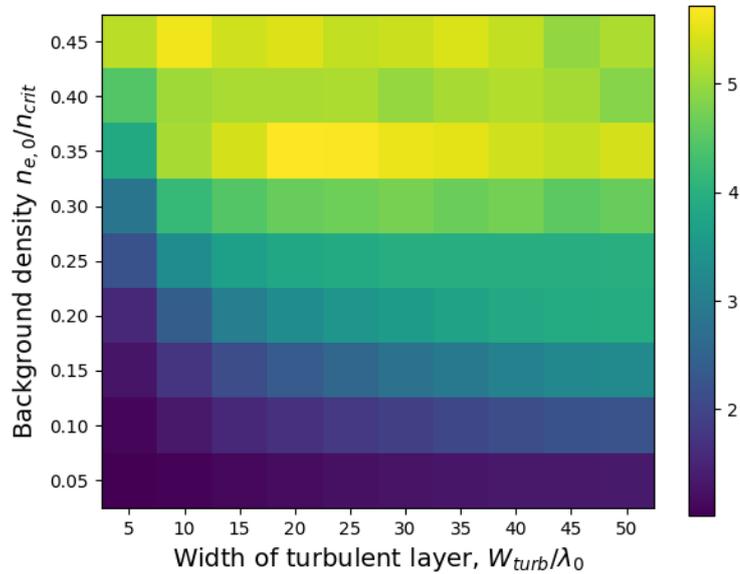


Figure 7.14: Dependence of broadening on width of the turbulence layer (W_{turb}) and background density ($n_{e,0}/n_{\text{crit}}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

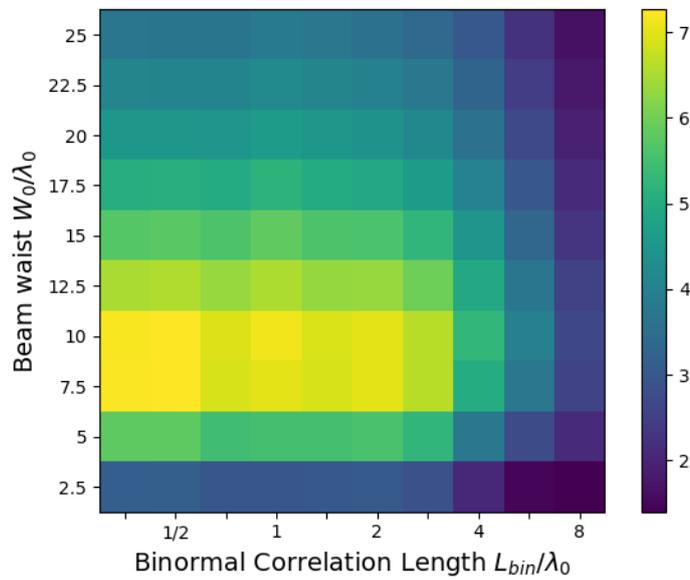


Figure 7.15: Dependence of broadening on microwave beam waist (w_0) and binormal correlation length (L_{bin}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

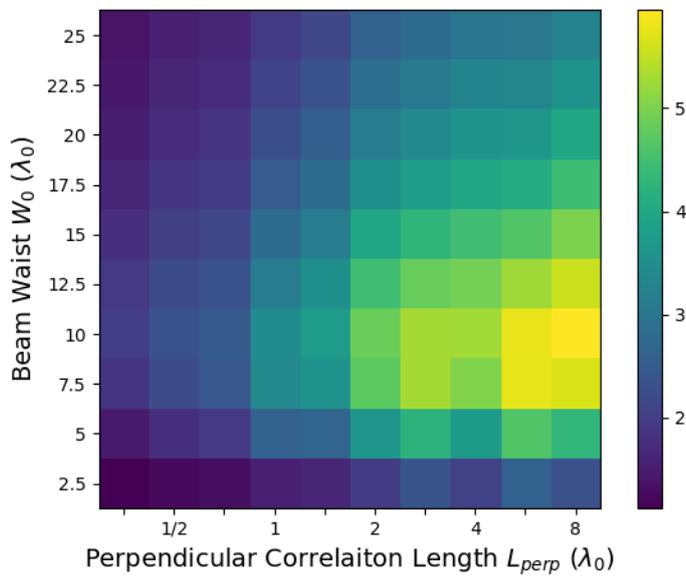


Figure 7.16: Dependence of broadening on microwave beam waist (w_0) and perpendicular correlation length (L_{perp}). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

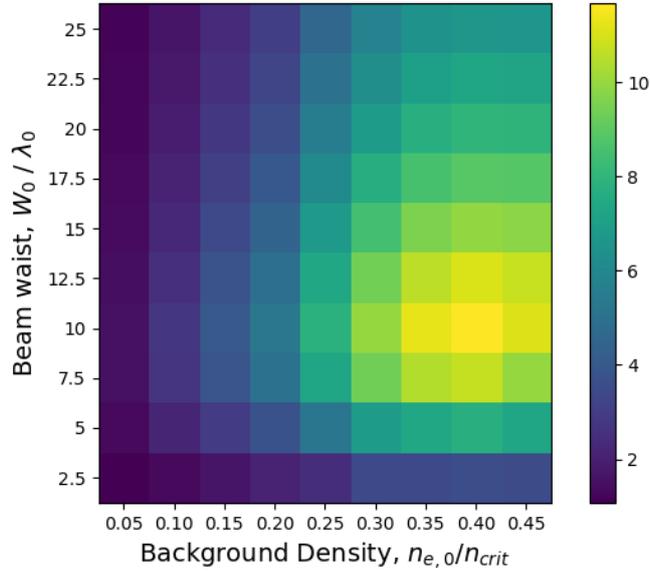


Figure 7.17: Dependence of broadening on microwave beam waist (w_0) and background density ($n_{e,0}/n_{crit}$). Colour represents broadening factor compared to a beam propagating through the same background density profile, with no turbulence present.

these extremal points, but only once the inseparability of the dependence is known, which it wasn't at the undertaking of these simulations.

The uncertainty in the broadening measurements from the simulations comes from the standard error on the ensemble average E_{RMS} . There is also an uncertainty associated with the standard error on the ensemble average correlation lengths, fluctuation level, and background density. This contributes to the error on the fit parameters, given by their covariance.

From Fig. 7.19, we can see that there is only a weak dependence on the sum of the logarithms of the correlation lengths and a stronger dependence on their difference. This corresponds to having a stronger dependence on the ratio $(L_{perp}/\lambda_0)^{0.4}/(L_{bin}/\lambda_0)^{0.9}$, meaning that having a larger L_{perp} compared to L_{bin} increases broadening. In fusion-relevant plasmas, the poloidal correlation length (which corresponds to L_{bin}) is usually longer than the radial correlation length (which corresponds to L_{perp}). This would result in less broadening. This rough trend matches the analytical result that (in the low fluctuation level limit) broadening is proportional to L_{perp}/L_{bin}^2 [90].

We can also see that the linear combination of $n_{e,0}/n_{crit}$ and $\delta n_e/n_{e,0}$ based on their difference has a weaker effect on the broadening than the linear combination based on their

sum. This suggests that perhaps a key consideration is the peak density of $n_{e,0}/n_{\text{crit}} + \delta n_e/n_{e,0}$ and how close it is to a cut-off. Due to the parameter ranges chosen, most points of the scan are outside of the limit where analytical theory can be applied, where we find the dependence of $n_{e,0}/n_{\text{crit}}$ and $\delta n_e/n_{e,0}$ is no longer separable and the broadening no longer scales with $(\delta n_e/n_{e,0})^2$ as predicted by theory [90].

The form of the dependence on microwave beam waist matches that found in [53], where, as w_0/λ_0 is increased, there is first a sharp increase in the broadening, followed by a gradual decrease. We believe this is due to the dependence of beam width on w_0/λ_0 even in the case where turbulence is not present, which can be seen in Eq. 3.31. The location of the peak is then dictated by the distance of the detection region (where the signal is measured) from the antenna.

Finally, the dependence of beam broadening on the width of the turbulence layer shows that as $W_{\text{turb}}/\lambda_0$ increases, so does broadening, until it reaches a maximal value and appears to level off. We believe this is due to a kind of saturation effect, where the beam can no longer become any broader without simply appearing as background noise. An example of this can be seen in Fig. (7.18) where the broadened ensemble average beam can be seen to be relatively flat and low amplitude. It could also be the cumulative affect of the previous broadening leading to a decrease in subsequent broadening because, as we know from the dependence on microwave beam waist, broader beams are broadened less by turbulence than narrower ones (once the non-monotonicity of beam width as a function of initial beam waist is accounted for).

The functional form of the dependence of broadening on plasma and beam parameters, based on the separability of parameters, is then

$$\begin{aligned} \Gamma = & 1 + C f_1 \left[\frac{(L_{\text{perp}}/\lambda_0)^{0.4}}{(L_{\text{bin}}/\lambda_0)^{0.9}} \right] f_2 \left[\left(\frac{L_{\text{perp}}}{\lambda_0} \right)^{0.9} \left(\frac{L_{\text{bin}}}{\lambda_0} \right)^{0.4} \right] \\ & \times f_3 \left[0.2 \left(\frac{n_{e,0}}{n_{\text{crit}}} \right) - 1.0 \left(\frac{\delta n_e}{n_{e,0}} \right) \right] f_4 \left[1.0 \left(\frac{n_{e,0}}{n_{\text{crit}}} \right) + 0.2 \left(\frac{\delta n_e}{n_{e,0}} \right) \right] \\ & \times f_5 \left[\frac{w_0}{\lambda_0} \right] f_6 \left[\frac{W_{\text{turb}}}{\lambda_0} \right] \end{aligned} \quad (7.6)$$

where each function is a vector of points from the point-wise fit, which can be interpolated between to find the broadening for a given set of parameters within the range investigated. The numerical values of these fit parameters can be found in Appendix A. The model presented

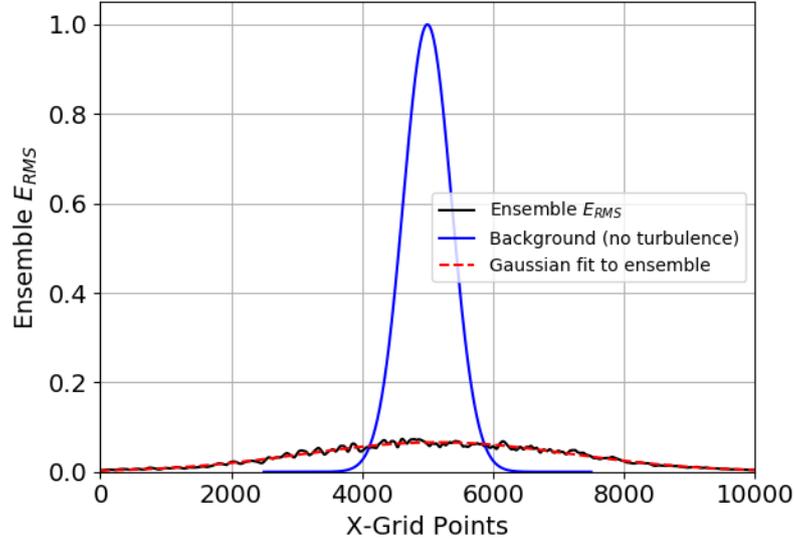


Figure 7.18: The results for a beam broadening simulation through an ensemble of turbulent profiles with all values as their base case apart from $W_{\text{turb}}/\lambda_0 = 50$. The Ensemble average E_{RMS} can be seen to be very low and flat such that it starts to look similar to background noise rather than a discernible Gaussian beam profile.

here is also available as a Python script [93].

7.4 Conclusions and Further Work

We used a 2D full-wave code to simulate microwave beams propagating through turbulent plasma density profiles to determine the dependence of the beam broadening on plasma and beam parameters. We carried out a series of 2D scans in order to identify which dependencies were separable. We found that the dependencies on $L_{\text{perp}}/\lambda_0$ and L_{bin}/λ_0 weren't separable from each other, and neither were the dependencies on $n_{e,0}/n_{\text{crit}}$ and $\delta n_e/n_{e,0}$. We found that all other dependencies were separable. We then found the dependence of each parameter, or linear combination of inseparable parameters, by means of a point-wise fit to the whole data set.

Where applicable, agreement with previous studies was found. However, it would be beneficial to compare this to other methods such as the analytical models [28,90], ray tracing simulations [27,70,91], and beam tracing simulations [29,92] in order to determine the regions of agreement and disagreement.

The inseparability of the dependence of broadening on the two turbulence correlation

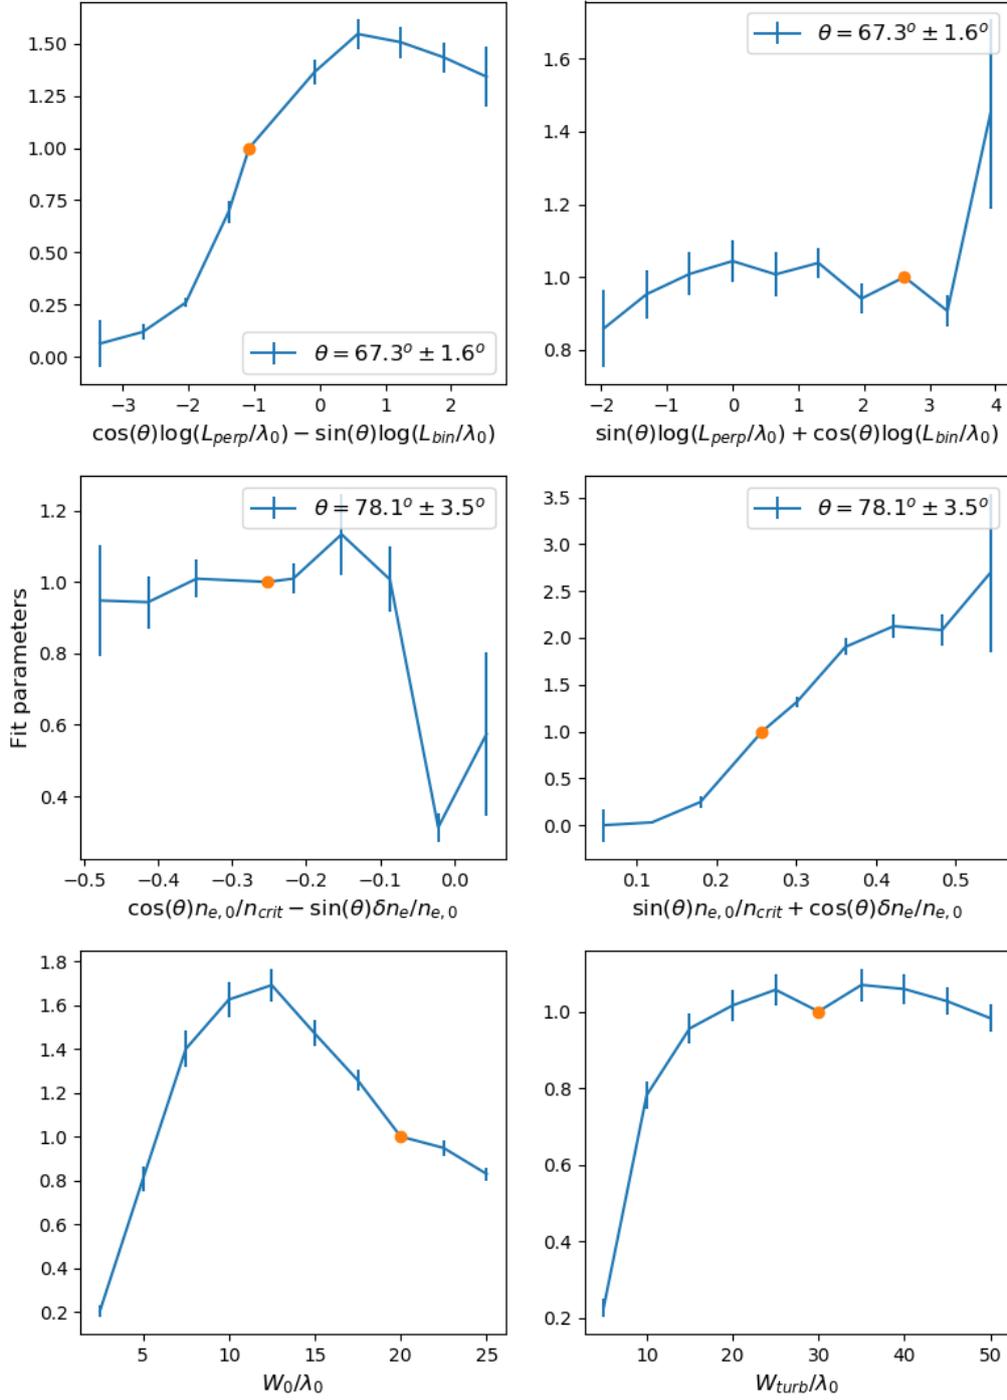


Figure 7.19: The normalised dependence of broadening on each parameter considered. For the case where the dependence of two parameters is not separable, the dependence on a linear combination of those parameters is shown instead. The orange points correspond to the base value of each parameter, which was used in scans where other parameters were varied. The fit parameters are normalised to this point. The solid lines serve as a guide to the eye.

lengths is an important result. Currently, when simulating beam broadening in fusion-relevant scenarios it is common only to have good data for the correlation length in one direction, resulting in both correlation lengths being set equal to each other. This has the potential to significantly under or over-predict the broadening of the microwave beam. The sensitivity of broadening to correlation length emphasises the need for diagnostic tools which can measure turbulence parameters to high degrees of precision. For example, from Fig. 7.5 the effect of uncertainty in correlation length measurement can be seen. If we take an example where only $L_{\text{bin}} = 2\lambda_0$ is measured, and for simulation purposes, the correlation lengths are set equal to each other despite $L_{\text{perp}} = 1\lambda_0$, the broadening would be over-predicted by a factor of 1.2 resulting in a large error on the predicted deposition profile.

A detailed understanding of the parametric dependence of beam broadening on plasma and beam parameters also introduces the possibility of being able to use beam broadening as a turbulence diagnostic. For example, if the other parameters are known well enough, the ratio of the turbulence correlation lengths could be deduced from the measured beam broadening.

Using the parameter dependencies provided here, it should now be possible to predict how much a beam will be broadened without further simulation, given the scenario falls within the parameter ranges investigated here. This can be done almost instantaneously, rather than the hours that would be required for an ensemble of full-wave simulations to be carried out. This ability to rapidly predict beam broadening could allow for its inclusion in integrated modelling and calculation during experiments, however, further work would be needed to extend the model's dimensionality if one wanted to simulate truly fusion-like scenarios. It is hoped that the work done here provides a solid framework for these next steps.

Chapter 8

Conclusions and Future Work

The main objective of this research project has been to better understand how microwaves propagate through fusion plasmas — crucially, how they propagate through magnetised plasmas where the plasma density varies on length scales comparable to the microwave wavelength.

In order to achieve this objective for parameters relevant to fusion reactors, a 2D full-wave code was needed. I decided to develop this myself from EMIT-3D (a cold plasma full-wave code utilising the FDTD method developed at the Universities of York and Strathclyde by previous PhD students). While other 2D full-wave codes already exist, I believed that the importance of having the absolute certainty that comes with fully knowing your own code, along with the flexibility of being able to edit it to suit your exact needs, outweighed the convenience of a pre-existing code.

The first portion of my PhD was therefore spent in code development followed by a thorough benchmark of the code. This led to a brief study of the OX mode conversion process in plasmas with steep density gradients. The objective then shifted to using EMIT-2D to further our understanding of the scattering effect density fluctuations have on microwave beams in fusion plasmas.

The first project was a study of ECRH beam broadening by plasma turbulence on the DIII-D tokamak. Three specific operating scenarios were selected, chosen to cover a wide range of parameters. Diagnostics on the device were used to build the plasma density profiles used in the simulations. The results of these simulations were then compared to an experimental measurement of the beam broadening. There were two main purposes of this work. One was to confirm that the additional broadening seen experimentally was predominantly

the result of turbulent scattering. The other was to provide an experimental benchmark of the code, confirming that it could make real-world predictions of this effect.

The second project was born out of the first. My work simulating scenarios for the DIII-D tokamak highlighted the sensitivity of beam broadening to plasma parameters. It also started to highlight the fact that the dependencies of beam broadening on various parameters might not be separable from each other. This led us to attempt to parameterise the beam broadening, considering pairs of parameters at a time in order to determine whether the two dependencies could be separated from each other, with the end goal being to develop an empirical formula or look-up table for beam broadening based on plasma and beam parameters. This work focused on X-mode propagation and the parameter ranges were chosen to encompass fusion relevant scenarios. The parameters of interest were background density, turbulence fluctuation level, turbulence correlation lengths (in both directions), width of the turbulent layer, and microwave beam waist.

The findings of these projects are summarised here, along with future plans for further work using EMIT-2D and EMIT-3D.

8.1 Conclusions

8.1.1 Code Development

The code development described in Chapter 4 was essential to the research carried out. EMIT-3D was further developed to improve the stability of the current density update equation and to make it suitable for simulation of X-mode propagation. The ability to launch an elliptically polarised beam was also added so that pure O-mode could be launched at an angle to the background magnetic field for the purpose of OX mode conversion studies.

EMIT-2D was developed from EMIT-3D to allow the simulation of much larger domains so as to be relevant for scenarios such as simulating a heating beam all the way from the antenna to where it will be absorbed.

Though the changes to EMIT-3D have not yet been used, they were propagated through to EMIT-2D and I'm sure they will be useful in future. The creation of EMIT-2D allowed for simulations that could fully compare to experiment and allowed for a wide ranging parameter scan to be carried out, which would not have been possible with the 3D code given constraints on computational time.

8.1.2 Benchmark and OX Mode Conversion Study

The benchmark described in Chapter 5 mostly found good agreement between EMIT-2D and two other full-wave codes, one of which used a similar FDTD approach and the other of which used a different Fourier method. Agreement was found for a range of scenarios of increasing complexity giving us good confidence in the new 2D code. Discrepancies between the codes were found close to the cut-off density when no magnetic field was present, but were found to be due to differing polarisations. In simulations where a magnetic field was present, the codes agreed comparatively poorly when between the O-mode cut-off and the UHR, but in this region, none of the codes are trying to accurately model the physical processes taking place as they are cold plasma codes. We therefore expect the exact values of the field to be highly dependent on numerical parameters and the method of power dissipation implemented. Outside of these regions, the three codes agreed well.

Once agreement between the codes was established, they were then used to carry out a study of OX mode conversion efficiency of relevance to MAST-U. The effect of density scale length was studied. The expected decrease in efficiency with increasing density scale length was reproduced. This is due to being more sensitive to mismatch on the optimum angle when the density gradient is less steep. However, an interesting additional effect was observed for short density scale lengths, where conversion efficiency dropped sharply for $k_0 L_n < 5$. This effect was observed in all three codes and was determined to be the result of the converted SX-mode tunnelling through the evanescent layer between the UHR and the R-wave cut-off to travel back out of the plasma along with the reflected O-mode. This is only possible when this evanescent layer is narrow enough to be comparable to the wavelength, which occurs for steep density gradients like those we might expect to see in H-mode operation on MAST-U, making it a potentially important consideration for those scenarios, particularly with regard to the path reflected power will take in the reactor and the damage it could cause to in-vessel components.

8.1.3 ECRH Beam Broadening on DIII-D

The research project described in Chapter 6 was able to confirm that the experimentally measured deposition profile of ECRH heating beams on DIII-D was broadened by turbulent scattering in the plasma edge. The simulation set-up was informed by diagnostic data from DIII-D for the background plasma density and turbulent density profile, and the beam was

initialised to match the heating beam on DIII-D.

The simulations were carried out from first principles and found good agreement with the experimental results. Given that the only possible source of broadening in the full-wave simulations was from scattering by turbulence, this suggests it must be the primary source of the broadening seen on DIII-D. Furthermore, though other parameters were changing as well, a clear scaling between fluctuation level and the broadening factor can be seen across the three scenarios. The fact that agreement was achieved for three operating scenarios covering a wide range of plasma parameters increases our confidence in this agreement. It also increases our confidence in the use of this approach to predict this effect in future tokamaks.

However, it should be noted that uncertainty on the diagnostic data that was used to construct the simulation density profiles resulted in significant uncertainty on the simulated beam broadening. For example, for the H-mode result of a broadening factor of 1.6 ± 0.5 we have an uncertainty of roughly 30%, but, given that the ‘zero-point’ of broadening is 1, it might perhaps be more accurate to describe this as an uncertainty of around 80%. This is most pronounced for the H-mode scenario due to the steep density gradients meaning small spatial shifts in measurements have a larger effect. This motivates the need for turbulence diagnostics with excellent spatial resolution as well as the need for accurate turbulence modelling to allow for the accurate prediction of this effect.

8.1.4 Parametric Dependence of Beam Broadening by Turbulence

The research project described in Chapter 7 was able to thoroughly characterise how beam broadening by electrostatic plasma turbulence depends on plasma and beam parameters. This was the first time such a study had been completed over parameter ranges encompassing fusion-relevant scenarios, where fluctuation level is high and structure size is comparable to microwave wavelength, and cut-off densities might appear in the turbulence. It was also the first time that the separability of dependencies on different parameters was investigated by conducting 2D parameter scans in pairwise combinations of parameters.

We found that the dependence on the two orthogonal turbulence correlation lengths could not be separated from each other, with the broadening strongly depending on the ratio $(L_{\text{perp}}/\lambda_0)^{0.4}/(L_{\text{bin}}/\lambda_0)^{0.9}$. We considered this to be an important result, as it shows that broadening will be less when the radial correlation length is shorter than the poloidal

correlation length (which is the case in fusion plasmas). It also motivates the need for good turbulence diagnostic data, as frequently one may only have experimental data for the correlation length in one direction at any particular location in the tokamak (as was the case for the previous study on DIII-D). In this case, it is common practice in simulations to set correlation lengths equal to each other, but this approach is liable to over-predict the broadening, potentially by quite large amounts.

We also found that dependences on background density and fluctuation level could not be separated from each other, with the broadening strongly depending on the sum $1.0(n_{e,0}/n_{\text{crit}}) + 0.2(\delta n_e/n_{e,0})$. We reason that this is because the key parameter is the peak density present, and more specifically, how close that peak density is to a cut-off.

All the other dependences were found to be separable.

The scaling of beam broadening with microwave beam waist was reproduced from previous studies and was determined to depend not only on the beam waist at the antenna but the distance between the antenna and the ‘receiver’. This is due to the fact that even if there was no turbulence present, measuring the width of the beam at the ‘receiver’ would reveal a non-monotonic dependence on the beam waist at the antenna due to beam divergence. The true dependence is that broader beams are broadened less by turbulence overall, but beams that start narrower at the antenna end up broader themselves even when turbulence isn’t present, so are also broadened less.

Beam broadening was found to increase with the width of the turbulence layer as would be expected, but after a certain point, a saturation effect was observed. We reasoned this could be due to the fact that broadening decreases with increasing beam width (as discussed above) so the broader a beam gets the less further broadening it experiences. It could also be due to reaching a point after which any further broadening would mean the beam starts to fade into background noise.

Where applicable, good agreement was found between the dependencies we observed and previous studies, including analytical theory.

Combining the results of all the parameter scans, we then came up with an empirical formula for the beam broadening, publishing a Python script for those who wish to use it to predict beam broadening in a matter of microseconds rather than the hours required for full-wave simulations.

8.2 Future Work

8.2.1 Beam Broadening by Turbulence

Whilst the work in this thesis has progressed our understanding of the broadening effect plasma turbulence has on microwave heating beams, there is still much work to be done.

One of the major outputs of the parameter scan work presented in Chapter 7 was a workflow that could be repeated or expanded. Applying this workflow to a specific machine could yield interesting and useful results. One of the main constraints with the original work was the sheer volume of simulations to be completed. Each new parameter or each additional degree of complexity increased the computational time required massively. Instead, constraining the parameter ranges with those relevant to a particular machine would reduce this cost, while allowing for more complex situations to be investigated. In addition, now that we know which dependencies are separable from each other, parameter scans in each pairwise combination of parameters are no longer required. This could allow for the simulations to include a full pedestal density profile along with the fluctuation level envelope, meaning pedestal width and steepness could also be investigated. It would also be of interest to investigate the effect of launching the beam at an angle to the background field, though to investigate the effect of toroidal injection angle would require the use of EMIT-3D which would vastly increase the computational cost. Developments in machine learning approaches could provide a helpful avenue for more easily and efficiently spanning a large parameter space.

Any and all of this work could also be repeated for O-mode polarisation, as we only considered X-mode. While the broadening effect is likely to be very similar, it would be worth confirming that to be the case.

Furthermore, in our work so far, we have only considered the broadening effect on the beam. We have not considered the scattering effect on the wave-vector components present in the beam and the impact this will have on absorption and heating and current drive efficiency. Whilst EMIT-2D and EMIT-3D are cold-plasma codes so cannot model the absorption of the wave themselves, they could be used to simulate the propagation through turbulence to the absorption region with the output used as input to codes which can model absorption and current drive. Given that absorption and current drive are sensitive to the wave vector, and we know that the scattering effect of turbulence will perturb the wave-vector make-up

of the beam, this might have a significant impact on the heating and current drive efficiency beyond the effect of broadening the deposition profile.

8.2.2 EMIT-3D Code Development

Whilst I spent some time developing the 3D version of the code, I did not end up using it for any of my research projects. In addition to the work described in Chapter 4, I also edited EMIT-3D to be truly 3D. Previously, all background parameters (plasma density and background magnetic field) were set to be uniform in the third dimension, with only wave quantities varying in all three dimensions. Whilst this is sufficient the vast majority of the time, a full 3D code might be necessary in cases where the curvature of the magnetic field in the toroidal direction is significant enough to become important.

In addition, I improved the parallelisation by making the code a hybrid of OpenMP and MPI, though the speed-up from this change has not yet been investigated. This also makes the file input and output system significantly easier to handle, as the input files only need to be opened once from each node rather than each core and only one output file is produced per node rather than one per core.

This development is still a work in progress, with some testing required before the code is used for research projects.

These changes were primarily made for the sake of the future research projects described here.

8.2.2.1 SAMI-2 Simulations

SAMI is a Synthetic Aperture Microwave Imaging diagnostic designed and built at the University of York [94] followed by the upgraded SAMI-2 [95].

SAMI-2 has multiple dual polarisation antennas that operate in the frequency range 20 – 40 GHz. The diagnostic has no focusing optics, instead using the phase difference between antennas to reconstruct a signal. It can operate in passive mode, detecting microwave emissions from the plasma to reconstruct a 2D image, or it can operate in active mode as a 2D Doppler back-scattering diagnostic.

In active mode, it can be used to calculate the magnetic pitch angle by finding the location of maximal blue and red shift in the reflected signal. This, in turn, can be used to reconstruct the edge current density profile which is crucial for understanding Edge Localised Modes.

A current PhD student at the University of York is in the process of installing SAMI-2 on MAST-U and analysing the data for this purpose.

Simulations will be key in the interpretation of diagnostic data. Plasma density and magnetic field profiles can be constructed for simulations using data from other diagnostics on MAST-U. A simulated SAMI-2 diagnostic can then be used to compare to experimental results to help determine what signal we would expect to detect. This will help determine whether the diagnostic is directly measuring the magnetic pitch angle, for example, or whether it is instead measuring a function of the magnetic pitch angle.

This will require some further code development, including two emitting antennas in EMIT-3D rather than the single antenna currently present. These antennas will need to be placed in the same physical locations that they appear on SAMI-2. It might also be beneficial to include the option to output data only at the receiving antennas on SAMI-2 (rather than across the whole antenna plane) for a true comparison to the diagnostic and to reduce data storage requirements.

8.2.2.2 OX Mode Conversion Simulations for MAST-U and STEP

The project to install gyrotrons on MAST-U for the purpose of EBW heating and current drive has already been discussed in Chapter 5. However, the work in that chapter focused solely on a slab geometry rather than a real-world tokamak. Simulation of the OX mode conversion process on MAST-U is ongoing. Furthermore, the UK prototype reactor STEP (Spherical Tokamak for Energy Production) is also considering scenarios with significant EBW heating and current drive contributions. This means understanding and being able to accurately predict the efficiency of this process will be very important.

While ray-tracing tools can be very useful for this, there are areas where full-wave codes are needed. For example, the reduction in OX mode conversion efficiency for steep density gradients due to the SX-mode being able to tunnel back out the plasma, as observed in Chapter 5, is an effect that can only be captured in a full-wave code. Another instance where full-wave codes such as EMIT-3D will likely be required is when considering the effect of edge turbulence on the mode conversion process.

Though the computational expense would be large, it would be beneficial to run full 3D simulations for a couple of pre-optimised cases on MAST-U and STEP to investigate the effect of edge turbulence on OX mode conversion efficiency. It would also be interesting to

investigate the effect on the XB mode conversion process, but this would require applying warm plasma corrections to EMIT-3D perhaps using a similar approach to Köhn *et al.* (2008) [96]. The results of these simulations could then be coupled to other tools such as ray tracers followed by Fokker-Planck codes to investigate the overall effect of edge turbulence on heating and current drive efficiency.

Appendix A

Parametric Dependence Fit Parameters

Here, we include the parameters from the fit as described in the Chapter 7. The results are presented in a Table (A.1). In order to find the predicted broadening, a value of each f as defined in Eq. (7.6) should be chosen. These are then multiplied together and multiplied by $C = 2.34 \pm 0.07$ before being added to 1 in order to find the factor by which the beam is broadened, such that

$$\Gamma = 1 + C f_1 f_2 f_3 f_4 f_5 f_6. \quad (\text{A.1})$$

Table A.1: Look up table for beam broadening by plasma turbulence based on plasma and beam parameters. p is the parameter (or combination of parameters where the dependence is not separable), f is the value of the fit function, and e is the error on that fit. $p_1 = (L_{\text{perp}}/\lambda_0)^{0.4}/(L_{\text{bin}}/\lambda_0)^{0.9}$, $p_2 = (L_{\text{perp}}/\lambda_0)^{0.9}(L_{\text{bin}}/\lambda_0)^{0.4}$, $p_3 = 0.2(n_{e,0}/n_{\text{crit}}) - 1.0(\delta n_e/n_{e,0})$, $p_4 = 1.0(n_{e,0}/n_{\text{crit}}) + 0.2(\delta n_e/n_{e,0})$, $p_5 = w_0/\lambda_0$, and $p_6 = W_{\text{turb}}/\lambda_0$

p_1	f_1	e_1	p_2	f_2	e_2	p_3	f_3	e_3	p_4	f_4	e_4	p_5	f_5	e_5	p_6	f_6	e_6
0.10	0.1	0.2	0.26	0.9	0.2	-0.48	0.9	0.2	0.06	0.0	0.2	2.5	0.20	0.03	5	0.23	0.03
0.15	0.12	0.04	0.40	0.95	0.07	-0.41	0.94	0.08	0.12	0.03	0.02	5.0	0.81	0.06	10	0.78	0.04
0.24	0.26	0.03	0.64	1.01	0.06	-0.35	1.01	0.06	0.18	0.25	0.07	7.5	1.40	0.09	15	0.96	0.04
0.38	0.70	0.06	1.00	1.04	0.06	-0.25	1	-	0.26	1	-	10.0	1.62	0.08	20	1.02	0.05
0.48	1	-	1.57	1.01	0.07	-0.22	1.01	0.05	0.30	1.32	0.06	12.5	1.69	0.08	25	1.06	0.05
0.95	1.36	0.07	2.48	1.04	0.05	-0.15	1.1	0.2	0.36	1.9	0.1	15.0	1.47	0.07	30	1	-
1.49	1.54	0.08	3.90	0.94	0.05	-0.09	1.0	0.1	0.42	2.1	0.2	17.5	1.25	0.05	35	1.07	0.05
2.35	1.51	0.08	6.13	1	-	-0.02	0.31	0.05	0.48	2.1	0.2	20.0	1	-	40	1.06	0.04
3.70	1.43	0.08	9.65	0.91	0.05	-0.04	0.6	0.3	0.54	2.7	0.9	22.5	0.95	0.04	45	1.03	0.04
5.82	1.3	0.2	15.2	1.4	0.3	-	-	-	-	-	-	25.0	0.83	0.03	50	0.98	0.04

List of References

- [1] Eurofusion. Tokamak configuration diagram. <https://www.eurofusion.org/>.
- [2] UKAEA. MAST Upgrade. <https://ccfe.ukaea.uk/research/mast-upgrade/>, 2023. Accessed: 2023-02-02.
- [3] T. Williams. *Full-wave simulation of high-frequency electromagnetic propagation through inhomogeneous plasma*. Ph.D. thesis, University of York, 2014.
- [4] Cineca HPC. MARCONI. <https://www.hpc.cineca.it/hardware/marconi>, 2023. Accessed: 2023-08-10.
- [5] World Nuclear Association. Nuclear fusion power. <https://world-nuclear.org/information-library/current-and-future-generation/nuclear-fusion-power.aspx>, 2023. Accessed: 2023-01-24.
- [6] J. D. Lawson. Some criteria for a power producing thermonuclear reactor. *Proceedings of the physical society. Section B*, 70(1):6, 1957.
- [7] C. Chen, J. Becker, and J. Farrell. Limits of energy confinement time and fusion energy gain in magnetic confinement fusion.
- [8] Y. M. Peng and D. J. Strickler. Features of spherical torus plasmas. *Nuclear Fusion*, 26(6):769, 1986.
- [9] D. Robinson. The physics of spherical confinement systems. *Plasma physics and controlled fusion*, 41(3A):A143, 1999.
- [10] A. Sykes, S. Team, N. Team, M. Team, et al. The spherical tokamak programme at culham. *Nuclear Fusion*, 39(9Y):1271, 1999.

-
- [11] G. D. Conway. Turbulence measurements in fusion plasmas. *Plasma Physics and Controlled Fusion*, 50(12):124026, 2008.
- [12] S. J. Zweben, J. A. Boedo, O. Grulke, C. Hidalgo, B. LaBombard, R. J. Maqueda, P. Scarin, and J. L. Terry. Edge turbulence measurements in toroidal fusion devices. *Plasma Physics and Controlled Fusion*, 49(7):S1, 2007.
- [13] W. Dorland, F. Jenko, M. Kotschenreuther, and B. N. Rogers. Electron temperature gradient turbulence. *Phys. Rev. Lett.*, 85:5579, 2000.
- [14] G. S. Lee and P. H. Diamond. Theory of ion-temperature-gradient-driven turbulence in tokamaks. *The Physics of Fluids*, 29(10):3291, 1986.
- [15] W. Horton. Drift waves and transport. *Rev. Mod. Phys.*, 71:735, 1999.
- [16] T. H. Stix. *Waves in plasmas*. Springer Science & Business Media, 1992.
- [17] R. J. La Haye. Neoclassical tearing modes and their control. *Physics of Plasmas*, 13(5):055501, 2006.
- [18] R. J. Buttery, et al. Neoclassical tearing modes. *Plasma Physics and Controlled Fusion*, 42(12B):B61, 2000.
- [19] R. Prater. Recent results on the application of electron cyclotron heating to tokamaks. *Journal of Fusion Energy*, 9(1):19, 1990.
- [20] T. Luce, Y. Lin-Liu, R. Harvey, G. Giruzzi, P. Politzer, B. Rice, J. Lohr, C. Petty, and R. Prater. Generation of localized noninductive current by electron cyclotron waves on the diii-d tokamak. *Physical review letters*, 83(22):4550, 1999.
- [21] Y. Ikeda, et al. Initial results of electron cyclotron range of frequency (ecrf) operation and experiments in jt-60u. *Fusion Engineering and Design*, 53(1):351, 2001.
- [22] R. Prater. Heating and current drive by electron cyclotron waves. *Physics of Plasmas*, 11(5):2349, 2004.
- [23] C. B. Forest, Y. S. Hwang, M. Ono, and D. S. Darrow. Internally generated currents in a small-aspect-ratio tokamak geometry. *Phys. Rev. Lett.*, 68:3559, 1992.

- [24] A. Ejiri, et al. Non-inductive plasma current start-up by EC and RF power in the TST-2 spherical tokamak. *Nuclear Fusion*, 49(6):065010, 2009.
- [25] H. P. Laqua. Electron bernstein wave heating and diagnostic. *Plasma Physics and Controlled Fusion*, 49(4):R1, 2007.
- [26] V. Shevchenko, M. O'Brien, D. Taylor, and A. S. and. Electron bernstein wave assisted plasma current start-up in MAST. *Nuclear Fusion*, 50(2):022004, 2010.
- [27] Y. Peysson, J. Decker, L. Morini, and S. Coda. RF current drive and plasma fluctuations. *Plasma Physics and Controlled Fusion*, 53(12):124028, 2011.
- [28] E. Sysoeva, F. da Silva, E. Gusakov, S. Heuraux, and A. Popov. Electron cyclotron resonance heating beam broadening in the edge turbulent plasma of fusion machines. *Nuclear Fusion*, 55(3):033016, 2015.
- [29] A. Köhn, L. Guidi, E. Holzhauser, O. Maj, E. Poli, A. Snicker, and H. Weber. Microwave beam broadening due to turbulent plasma density fluctuations within the limit of the born approximation and beyond. *Plasma Physics and Controlled Fusion*, 60(7):075006, 2018.
- [30] H. J. Hartfuss. RF techniques in plasma diagnostics. *Plasma Physics and Controlled Fusion*, 40(8A):A231, 1998.
- [31] N. C. L. Jr., H. Bindslev, H. Park, J. Sánchez, G. Taylor, and C. X. Yu. Chapter 3: Microwave diagnostics. *Fusion Science and Technology*, 53(2):335, 2008.
- [32] R. Fitzpatrick. *Plasma physics: an introduction*. Crc Press, 2014.
- [33] M. G. Senstius, S. K. Nielsen, R. Vann, and S. K. Hansen. Particle-in-cell simulations of parametric decay instabilities at the upper hybrid layer of fusion plasmas to determine their primary threshold. *Plasma Physics and Controlled Fusion*, 62(2):025010, 2019.
- [34] S. Hansen, S. Nielsen, J. Stober, J. Rasmussen, M. Salewski, M. Stejner, A. U. Team, et al. Power threshold and saturation of parametric decay instabilities near the upper hybrid resonance in plasmas. *Physics of Plasmas*, 26(6), 2019.
- [35] E. V. Appleton. Wireless studies of the ionosphere. *Institution of Electrical Engineers- Proceedings of the Wireless Section of the Institution*, 7(21):257, 1932.

- [36] D. R. Hartree. The propagation of electromagnetic waves in a refracting medium in a magnetic field. *Mathematical Proceedings of the Cambridge Philosophical Society*, 27(1):143–162, 1931.
- [37] E. R. Tracy, A. J. Brizard, A. S. Richardson, and A. N. Kaufman. *Ray Tracing and Beyond: Phase Space Methods in Plasma Wave Theory*. Cambridge University Press, 2014.
- [38] E. Poli, G. Pereverzev, A. Peeters, and M. Bornatici. Ec beam tracing in fusion plasmas. *Fusion Engineering and Design*, 53(1):9 , 2001.
- [39] D. B. Davidson. *Computational Electromagnetics for RF and Microwave Engineering*. Cambridge University Press, 2nd edition, 2010.
- [40] B. F. Shorr. *The wave finite element method*. Springer Science & Business Media, 2012.
- [41] A. Taflove and S. C. Hagness. *Computational electrodynamics: the finite-difference time-domain method*. Artech house, 2005.
- [42] J. B. Schneider. Understanding the finite-difference time-domain method. *School of electrical engineering and computer science Washington State University*, 2010.
- [43] R. Fitzpatrick. Introduction to computational physics. *Retrieved*, 4(20):2012, 2011.
- [44] D. Tskhakaya. *The particle-in-cell method*. Springer, 2008.
- [45] J. M. Dawson. Particle simulation of plasmas. *Rev. Mod. Phys.*, 55:403, 1983.
- [46] J. Decker. *Electron Bernstein Wave Current Drive Modeling in Toroidal Plasma Confinement*. Ph.D. thesis, 2005.
- [47] M. Thomas. *3D full-wave modelling of microwave interactions with plasma density fluctuations*. Ph.D. thesis, University of York, 2018.
- [48] D. Woodward. *Microwave propagation and interaction in fusion plasma*. Ph.D. thesis, University of Strathclyde, 2022.
- [49] K. Yee. Numerical solution of initial boundary value problems involving maxwell’s equations in isotropic media. *IEEE Transactions on antennas and propagation*, 14(3):302, 1966.

- [50] J. Irby, S. Horne, I. Hutchinson, and P. Stek. 2d full-wave simulation of ordinary mode reflectometry. *Plasma physics and controlled fusion*, 35(5):601, 1993.
- [51] H. Hutchinson. *Principles of Plasma Diagnostics*. Cambridge University Press, 1987.
- [52] F. Hansen, J. P. Lynov, and P. Michelsen. The oxb mode conversion scheme for ecrh of a high-density tokamak plasma. *Plasma physics and controlled fusion*, 27(10):1077, 1985.
- [53] A. Köhn, E. Holzhauser, J. Leddy, M. B. Thomas, and R. G. L. Vann. Influence of plasma turbulence on microwave propagation. *Plasma Physics and Controlled Fusion*, 58(10):105008, 2016.
- [54] A. Köhn-Seemann, B. E. Eliasson, S. J. Freethy, L. A. Holland, and R. G. Vann. Benchmarking full-wave codes for studying the o-sx mode conversion in mast upgrade. In *EPJ Web of Conferences*, volume 277, page 01010. EDP Sciences, 2023.
- [55] CCFE. Mast upgrade research plan. https://ccfe.ukaea.uk/wp-content/uploads/2019/12/MAST-U_RP_2019_v1.pdf, 2019.
- [56] A. Köhn. *Investigation of microwave heating scenarios in the magnetically confined low-temperature plasma of the stellarator TJ-K*. Ph.D. thesis, Universität Stuttgart, 2010.
- [57] B. Eliasson, K. Ronald, A. W. Cross, D. C. Speirs, and D. Woodward. Step report, wp 4.2(a): Fourier full-wave (ffw) simulations of o-x-b mode conversion. Technical report, UKAEA, 2020.
- [58] A. Köhn, J. Jacquot, M. W. Bongard, S. Gallian, E. T. Hinson, and F. A. Volpe. Full-wave modeling of the o-x mode conversion in the pegasus toroidal experiment. *Physics of Plasmas*, 18(8), 2011.
- [59] M. W. Brookman, et al. Broadening of microwave heating beams in the diii-d tokamak by edge turbulence. *Nuclear Fusion*, 63(4):044001, 2023.
- [60] H.-J. Hartfuss and T. Geist. *Fusion Plasma Diagnostics with mm-waves: an introduction*. John Wiley & Sons, 2013.

-
- [61] M. Bornatici, R. Cano, O. De Barbieri, and F. Engelmann. Electron cyclotron emission and absorption in fusion plasmas. *Nuclear Fusion*, 23(9):1153, 1983.
- [62] V. Erckmann and U. Gasparino. Electron cyclotron resonance heating and current drive in toroidal fusion plasmas. *Plasma physics and controlled fusion*, 36(12):1869, 1994.
- [63] M. Shimada, et al. Overview and summary. *Nuclear Fusion*, 47(6):S1, 2007.
- [64] H. Zohm. Stabilization of neoclassical tearing modes by electron cyclotron current drive. *Physics of plasmas*, 4(9):3433, 1997.
- [65] C. C. Hegna and J. D. Callen. On the stabilization of neoclassical magnetohydrodynamic tearing modes using localized current drive or heating. *Physics of Plasmas*, 4(8):2940, 1997.
- [66] O. Sauter, M. A. Henderson, G. Ramponi, H. Zohm, and C. Zucca. On the requirements to control neoclassical tearing modes in burning plasmas. *Plasma Physics and Controlled Fusion*, 52(2):025002, 2010.
- [67] E. Ott, B. Hui, and K. R. Chu. Theory of electron cyclotron resonance heating of tokamak plasmas. *The Physics of Fluids*, 23(5):1031, 1980.
- [68] B. Hui, E. Ott, P. Bonoli, and P. Guzdar. Scattering of electron cyclotron resonance heating waves by density fluctuations in tokamak plasmas. *Nuclear Fusion*, 21(3):339, 1981.
- [69] F. Hansen, J. Lynov, P. Michelsen, and H. Pécseli. Ordinary wave propagation in a tokamak with random density fluctuations. *Nuclear Fusion*, 28(5):769, 1988.
- [70] C. Tsironis, A. G. Peeters, H. Isliker, D. Strintzi, I. Chatziantonaki, and L. Vlahos. Electron-cyclotron wave scattering by edge density fluctuations in ITER. *Physics of Plasmas*, 16(11):112510, 2009.
- [71] E. Poli, et al. On recent results in the modelling of neoclassical-tearing-mode stabilization via electron cyclotron current drive and their impact on the design of the upper ec launcher for iter. *Nuclear Fusion*, 55(1):013023, 2015.

- [72] T. Williams, A. Köhn, M. O'Brien, and R. G. Vann. Propagation in 3d of microwaves through density perturbations. *Plasma Physics and Controlled Fusion*, 56(7):075010, 2014.
- [73] Y. Peysson, J. Decker, and L. Morini. A versatile ray-tracing code for studying rf wave propagation in toroidal magnetized plasmas. *Plasma Physics and Controlled Fusion*, 54(4):045003, 2012.
- [74] A. Snicker, E. Poli, O. Maj, L. Guidi, A. Köhn, H. Weber, G. Conway, M. Henderson, and G. Saibene. The effect of density fluctuations on electron cyclotron beam broadening and implications for iter. *Nuclear Fusion*, 58(1):016002, 2017.
- [75] O. Chellaï, et al. Millimeter-wave beam scattering by field-aligned blobs in simple magnetized toroidal plasmas. *Physical review letters*, 120(10):105001, 2018.
- [76] O. Chellaï, et al. Millimeter-wave beam scattering and induced broadening by plasma turbulence in the tcv tokamak. *nuclear fusion*, 61(6):066011, 2021.
- [77] M. Brookman, et al. Resolving ecrh deposition broadening due to edge turbulence in diii-d. *Physics of Plasmas*, 28(4), 2021.
- [78] W. S. for The DIII-D Team. Diii-d research advancing the scientific basis for burning plasmas and fusion energy. *Nuclear Fusion*, 57(10):102018, 2017.
- [79] J. Lohr, et al. The electron cyclotron resonant heating system on the diii-d tokamak. *Fusion science and technology*, 48(2):1226, 2005.
- [80] E. Kolemen, A. Welander, R. La Haye, N. Eidiētis, D. Humphreys, J. Lohr, V. Noraky, B. Penafior, R. Prater, and F. Turco. State-of-the-art neoclassical tearing mode control in diii-d using real-time steerable electron cyclotron current drive launchers. *Nuclear Fusion*, 54(7):073020, 2014.
- [81] T. Luce, J. DeBoo, N. Howard, C. Liu, and C. Petty. Search for threshold behavior in diii-d electron heat transport. In *Proc. 32nd EPS Conf. on Plasma Phys. ECA (27 June, 2005 Tarragona, Spain)*, volume 29, page P5. 2005.
- [82] J. H. Slief, R. van Kampen, M. W. Brookman, and M. van Berkel. Extension of the flux fit method for estimating power deposition profiles. *Physics of Plasmas*, 29(1), 2022.

- [83] R. Prater, et al. Benchmarking of codes for electron cyclotron heating and electron cyclotron current drive under iter conditions. *Nuclear Fusion*, 48(3):035006, 2008.
- [84] K. Kirov, F. Leuterer, G. Pereverzev, F. Ryter, W. Suttrop, A. U. Team, et al. Ecrh power deposition studies in asdex upgrade. *Plasma physics and controlled fusion*, 44(12):2583, 2002.
- [85] F. Leuterer, K. Kirov, G. Pereverzev, F. Ryter, and D. Wagner. Modulated ecrh power deposition in asdex upgrade. *Nuclear fusion*, 43(8):744, 2003.
- [86] T. Rhodes, J.-N. Leboeuf, R. Sydora, R. Groebner, E. Doyle, G. McKee, W. Peebles, C. Rettig, L. Zeng, and G. Wang. Comparison of turbulence measurements from diii-d low-mode and high-performance plasmas to turbulence simulations and models. *Physics of Plasmas*, 9(5):2141, 2002.
- [87] B. D. Dudson and J. Leddy. Hermes: global plasma edge fluid turbulence simulations. *Plasma Physics and Controlled Fusion*, 59(5):054010, 2017.
- [88] L. Holland, A. Köhn-Seemann, and R. Vann. Parametric dependence of microwave beam broadening by plasma density turbulence. *Nuclear Fusion*, 63(5):056013, 2023.
- [89] H. Zohm, et al. Experiments on neoclassical tearing mode stabilization by ECCD in ASDEX Upgrade. *Nuclear Fusion*, 39(5):577, 1999.
- [90] P. Tretinnikov, E. Gusakov, and S. Heuraux. X-mode beam broadening in turbulent plasma. *Plasma Physics and Controlled Fusion*, 63(8):085003, 2021.
- [91] N. Bertelli, A. A. Balakin, E. Westerhof, O. E. Garcia, A. H. Nielsen, and V. Naulin. The influence of the edge density fluctuations on electron cyclotron wave beam propagation in tokamaks. *Journal of Physics: Conference Series*, 260:012002, 2010.
- [92] A. Snicker, L. Guidi, A. Köhn, O. Maj, H. Weber, and E. Poli. The effect of density fluctuations on ECRH beam broadening and implications to NTM mitigation on ITER. In *APS Division of Plasma Physics Meeting Abstracts*, volume 2016, pages PO6–005. 2016.
- [93] L. Holland. Beam broadening calculator: a python script. <https://doi.org/10.15124/a0a0106c-42e4-47a0-b5ae9fc5171fcc53>.

- [94] S. Freethy. *Synthetic aperture imaging of B-X-O mode conversion*. Ph.D. thesis, University of York, 2012.
- [95] J. Allen. *Design of the Synthetic Aperture Microwave Imager-2 for measurement of the edge current density on MAST-U*. Ph.D. thesis, University of York, 2021.
- [96] A. Köhn, A. Cappa, E. Holzauer, F. Castejón, Á. Fernández, and U. Stroth. Full-wave calculation of the O–X–B mode conversion of gaussian beams in a cylindrical plasma. *Plasma Physics and Controlled Fusion*, 50(8):085018, 2008.