

Library Declaration and Deposit Agreement

1. STUDENT DETAILS

Benjamin John Chapman
1556902

2. THESIS DEPOSIT

- 2.1 I understand that under my registration at the University, I am required to deposit my thesis with the University in BOTH hard copy and in digital format. The digital version should normally be saved as a single pdf file.
- 2.2 The hard copy will be housed in the University Library. The digital version will be deposited in the University's Institutional Repository (WRAP). Unless otherwise indicated (see 2.3 below) this will be made openly accessible on the Internet and will be supplied to the British Library to be made available online via its Electronic Theses Online Service (EThOS) service. [At present, theses submitted for a Masters degree by Research (MA, MSc, LLM, MS or MMedSci) are not being deposited in WRAP and not being made available via EThOS. This may change in future.]
- 2.3 In exceptional circumstances, the Chair of the Board of Graduate Studies may grant permission for an embargo to be placed on public access to the hard copy thesis for a limited period. It is also possible to apply separately for an embargo on the digital version. (Further information is available in the Guide to Examinations for Higher Degrees by Research.)
- 2.4 (a) Hard Copy I hereby deposit a hard copy of my thesis in the University Library to be made publicly available to readers immediately.
I agree that my thesis may be photocopied.
- (b) Digital Copy I hereby deposit a digital copy of my thesis to be held in WRAP and made available via EThOS.
My thesis can be made publicly available online.

3. GRANTING OF NON-EXCLUSIVE RIGHTS

Whether I deposit my Work personally or through an assistant or other agent, I agree to the following: Rights granted to the University of Warwick and the British Library and the user of the thesis through this agreement are non-exclusive. I retain all rights in the thesis in its present version or future versions. I agree that the institutional repository administrators and the British Library or their agents may, without changing content, digitise and migrate the thesis to any medium or format for the purpose of future preservation and accessibility.

4. DECLARATIONS

(a) I DECLARE THAT:

- I am the author and owner of the copyright in the thesis and/or I have the authority of the authors and owners of the copyright in the thesis to make this agreement. Reproduction of any part of this thesis for teaching or in academic or other forms of publication is subject to the normal limitations on the use of copyrighted materials and to the proper and full acknowledgement of its source.
- The digital version of the thesis I am supplying is the same version as the final, hardbound copy submitted in completion of my degree, once any minor corrections have been completed.
- I have exercised reasonable care to ensure that the thesis is original, and does not to the best of my knowledge break any UK law or other Intellectual Property Right, or contain any confidential material.
- I understand that, through the medium of the Internet, files will be available to automated agents, and may be searched and copied by, for example, text mining and plagiarism detection software.

(b) IF I HAVE AGREED (in Section 2 above) TO MAKE MY THESIS PUBLICLY AVAILABLE DIGITALLY, I ALSO DECLARE THAT:


- I grant the University of Warwick and the British Library a licence to make available on the Internet the thesis in digitised format through the Institutional Repository and through the British Library via the EThOS service.
- If my thesis does include any substantial subsidiary material owned by third-party copyright holders, I have sought and obtained permission to include it in any version of my thesis available in digital format and that this permission encompasses the rights that I have granted to the University of Warwick and to the British Library.

5. LEGAL INFRINGEMENTS

I understand that neither the University of Warwick nor the British Library have any obligation to take legal action on behalf of myself, or other rights holders, in the event of infringement of intellectual property rights, breach of contract or of any other right, in the thesis.

Please sign this agreement and return it to the Graduate School Office when you submit your thesis.

Student's signature:Date:



**Ion cyclotron emission from energetic ion
populations in fusion plasmas**

by

Benjamin John Chapman

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Physics

March 2019

THE UNIVERSITY OF
WARWICK

Contents

List of Tables	iv
List of Figures	v
Acknowledgments	xxvi
Declarations	xxvii
Abstract	xxix
Chapter 1 Introduction	1
1.1 Nuclear Fusion	1
1.2 Plasma physics	4
1.2.1 Single particle effects	4
1.2.2 Plasma waves	6
1.3 Tokamak plasmas	12
1.3.1 Confinement	12
1.3.2 Fast particles	14
1.3.3 Edge Localized Modes (ELMS)	17
1.4 Ion cyclotron emission and the Magnetoacoustic cyclotron instability	18
1.4.1 Ion cyclotron emission (ICE)	19
1.4.2 The Magnetoacoustic Ion cyclotron instability (MCI)	24
Chapter 2 Computational methods: Particle-In-Cell codes and higher order spectra	29
2.1 Particle-In-Cell codes	29
2.1.1 Field update	31
2.1.2 Particle Push	32
2.1.3 Particle shape functions	34
2.1.4 Practical and computational considerations	35

2.1.5	Cold plasma dispersion relation	38
2.2	Bicoherence and Bispectral analysis	39
Chapter 3 Simulations of the MCI - Resolution testing and Parameter scans		43
3.1	Domain size	44
3.2	Cell size	46
3.3	Summary	50
Chapter 4 Simulations of ICE in KSTAR		52
4.1	Fusion born proton frequency chirping	52
4.1.1	ELMs and ICE in KSTAR	52
4.1.2	Single particle orbits	54
4.1.3	Self consistent simulations of the MCI	54
4.1.4	ICE intensity as a function of frequency and electron number density	57
4.1.5	Mapping between computational results and KSTAR observations	57
4.1.6	Upward chirping	60
4.2	The effect of the lower hybrid frequency	62
4.3	Nonlinear wave-wave interactions	67
4.3.1	Bicoherence analysis	69
4.3.2	Density dependence of downward chirping	73
4.4	Steady state deuterium NBI ICE	74
4.4.1	Identifying NBI deuterons in KSTAR that could relax via the MCI	77
4.4.2	Comparison between kinetic simulations and experiment	78
4.5	Conclusions	85
Chapter 5 Simulations of ICE in JET and ASDEX Upgrade		89
5.1	JET	89
5.2	ASDEX Upgrade	96
5.3	Conclusions	109
Chapter 6 Simulations of core ICE		111
6.1	Description of simulations	112
6.2	Linear physics	113
6.3	Nonlinear physics	125

6.4	Conclusions	137
Chapter 7	Simulations with multiple minority ion species	140
7.1	Simulations of KSTAR ICE with NBI deuterons and fusion-born protons	140
7.2	Preliminary simulations of helium ash pumping in JET core plasmas	149
7.3	Conclusions	153
Chapter 8	Summary	156
8.1	KSTAR ICE	156
8.2	JET and AUG ICE	159
8.3	Core ICE	160
8.4	Preliminary simulations of helium ash pumping in JET core plasmas	161
Appendix A	Hybrid version of EPOCH	162
A.1	Introduction to hybrid codes	162
A.2	Hybrid field equations	163
A.3	Numerical implementation	164
	A.3.1 Additions to the EPOCH input deck	167
A.4	Summary and future	167
Appendix B	Mapping between experimental and simulation harmonics	168

List of Tables

1.1	Waves in cold uniform plasmas.	9
1.2	Waves in warm uniform plasmas.	11
1.3	Timeline of experimental ICE observations in MCF plasmas.	23

List of Figures

1.1	A plasma trapped between magnetic mirrors [Hutchinson, 2001].	2
1.2	Cold plasma dispersion relation for parallel propagating (red) and perpendicularly propagating (blue) waves. Only waves in the ion cyclotron range of frequencies are shown. The y-axis is normalised to the deuteron cyclotron frequency, while the x-axis is normalised to the deuteron cyclotron frequency divided by the Alfvén speed. The horizontal dashed line denotes the deuteron cyclotron frequency, the diagonal dashed line denotes the Alfvén speed, and the horizontal solid line denotes the lower hybrid frequency for perpendicular propagation.	7
1.3	The electromagnetic coil configuration of a conventional tokamak showing the resulting magnetic field [EuroFusion].	13
1.4	Poloidal cross section showing the cylindrical coordinate system used to model a tokamak. [Freidberg, 1987; Sauter and Medvedev, 2013]	15
1.5	Left: Trapped orbit of a 3.5MeV α particle in a JET-like equilibrium. Right: Passing orbit of a 3.5MeV α particle in a JET-like equilibrium.	17
1.6	Left: ICE intensity as a function of frequency for fusion born ICE in JET. Reproduced from [Cottrell et al., 1993]. Right: ICE intensity as a function of frequency for fusion born ICE in TFTR. Reproduced from [Cauffman et al., 1995]. In both plots, evenly spaced sequential harmonics of the alpha particle cyclotron frequency can be seen.	20
1.7	Spectrogram of E_y component of the field measurement by the Van Allen Probes spacecraft. Reproduced from [Posch et al., 2015].	24
1.8	Phase space plots showing the variants of the distribution function used to represent the minority ion species in MCI simulations. All distributions have identical values of the initial perpendicular velocity u_0 . The parallel drift velocity v_d , parallel spread v_r , and perpendicular spread u_r are expressed as multiples of u_0 . Top left: Corresponds to Eq. 1.22. Top right: Corresponds to Eq. 1.23. Bottom left: Corresponds to Eq. 1.24. Bottom right: Corresponds to Eq. 1.25.	28

2.1	Schematic of the 1D Yee grid used in the 1D3V version of the EPOCH PIC code.	32
2.2	Schematic of the first order tophat (blue) and second order triangular (red) shape functions used in the EPOCH PIC code. This figure has been reproduced from page 21 of the most recent EPOCH developers manual (at the time of writing).	35
2.3	Warm plasma dispersion relation as simulated by EPOCH. Left: Perpendicularly propagating waves. Right: Parallel propagating waves. Shading indicates the spectral power in B_z field component and is plotted on a log10 scale.	39
3.1	Change in energy density of particles and electromagnetic fields as a function of time. Time is plotted in units of the alpha particle gyro period. The initial energy densities of the E_x , B_y , and ΔB_z field component are all zero; while the initial energy densities of the electrons, thermal deuterons, and NBI deuterons are $\sim 2.4 \times 10^3 \text{Jm}^{-3}$, $\sim 2.4 \times 10^3 \text{Jm}^{-3}$, and $\sim 5.6 \times 10^3 \text{Jm}^{-3}$ respectively.	45
3.2	Frequency-wavenumber space of the MCI. Shading indicates the spectral power in B_z field component and is plotted on a log10 scale.	46
3.3	Spectral intensity of the B_z field component of the MCI simulations. Power is obtained by integrating between $0[\omega_{c\alpha}/V_A] < k < 35[\omega_{c\alpha}/V_A]$.	46
3.4	As Fig. 3.1. Panels a), b) and c) correspond to simulations using triangular simulations, panels d), e) and f) correspond to simulations using higher order spline shape functions. Row one, panels a) and d): $\Delta x \sim \lambda_D$. Row two, panels b) and e): $\Delta x \sim 3\lambda_D$. Row three, panels c) and f): $\Delta x \sim 5\lambda_D$.	48
3.5	Spectral intensity of the B_z field component of the MCI simulations. Subplots are the same as in Fig. 3.4.	49
3.6	Percentage change in total simulation energy as a function of time. Time is plotted in units of the alpha particle gyro period. The applied magnetic field $ \mathbf{B} = 2.1\text{T}$ has been excluded from the calculation. Subplots are the same as in Figs. 3.4 and 3.5.	50
4.1	Temporal evolution of ICE amplitude (upper plot) and spectrum (lower plot) during an ELM crash in KSTAR plasma 11513. Time is measured relative to the moment chirping bursts are observed to begin during the ELM crash ($\sim 100\mu\text{s}$ after the start of the crash). The horizontal dashed lines in the spectrogram indicate proton cyclotron harmonics.	53

4.2	Poloidal projection of 3.0 MeV fusion proton orbits in the model KSTAR equilibrium with initial velocity vectors slightly offset from the co-current toroidal direction. The blue boxes in both the left and right panels designate the emitting region for the ICE and are centred on the radial location that corresponds to the proton cyclotron frequency whose harmonics are excited in the chirping ICE spectrum. The horizontal extent of these boxes corresponds to the local Larmor radius of cyclotron gyration of an emitting proton.	55
4.3	Top panel: experimentally-measured fast RF burst spectrogram from KSTAR plasma 11462 with $B_0 = 1.7\text{T}$ and average electron number density before the ELM crash $\langle n_e \rangle = 2.5 \times 10^{19}\text{m}^{-3}$. Downward step-wise frequency chirping with proton cyclotron frequency $f_{cp} \sim 21.5\text{MHz}$ is apparent. Lower panels: frequency versus electron number density plots for the nonlinear stage of MCI simulations where $B_z = B_{cyc} \approx 1.41\text{T}$ has been inferred from the data in the top panel. Shading indicates the \log_{10} of the spectral power in the fluctuating part of the B_z field component of each simulation.	58
4.4	Top panel: experimentally-measured fast RF burst spectrogram from KSTAR plasma 11513 with $B_0 = 1.99\text{T}$ and average electron number density before the ELM crash $\langle n_e \rangle = 2.6 \times 10^{19}\text{m}^{-3}$. Downward step-wise frequency chirping with $f_{cp} \sim 25\text{MHz}$ is apparent. Lower panels: as the lower panel of Fig. 4.3 but with $B_z = B_{cyc} \approx 1.64\text{T}$	59
4.5	Local electron number density during an ELM crash in KSTAR inferred from downward chirping ICE measurements compared with saturated MCI simulations. Left (right): fitted and estimated electron number density and time values corresponding to KSTAR plasma 11462 (plasma 11513). The green fit suggests that local electron number density declines approximately linearly with time during the early stages of the ELM crash.	61
4.6	Left panel: experimentally-measured fast RF burst spectrogram from KSTAR plasma 11474 with $B_0 = 2.27\text{T}$ and $\langle n_e \rangle = 2.5 \times 10^{19}\text{m}^{-3}$. Upward step-wise frequency chirping with $f_{cp} \sim 28\text{MHz}$ is apparent. Right panel: as the lower panel of Fig. 4.3 but with $B_z = B_{cyc} \approx 1.84\text{T}$	61

4.7	Rising electron number density during an ELM crash inferred from upward chirping ICE measurements in KSTAR plasma 11474 combined with saturated MCI simulations at different densities, see Fig. 4.6. The green fit suggests electron number density rises approximately linearly with time during the early stages of the ELM crash.	62
4.8	Left: The lower hybrid frequency ω_{LH} as a function of electron particle density n_e for the KSTAR plasma parameters used in the PIC simulations. ω_{LH} is normalised to the proton cyclotron frequency ω_{cp} while n_e is normalised to 10^{19} . Each line denotes ω_{LH} for different values of central magnetic field B_0 spanning 1.7T at the top (blue) to 2.3T at the bottom (black) in steps of 0.1T. Right: ω_{LH} as a function of B_0 calculated for different values of n_e spanning $2.4 \times 10^{19} \text{m}^{-3}$ at the top (blue) to $0.2 \times 10^{19} \text{m}^{-3}$ at the bottom (magenta) in steps of $0.2 \times 10^{19} \text{m}^{-3}$.	63
4.9	Spectral power as a function of plasma density obtained from multiple PIC simulations for several values of magnetic field strength corresponding to the nonlinear saturated phase of the MCI. Shading indicates the \log_{10} of the spectral power in the B_z field component of each simulation. All panels comprise a series of simulations, each contributing a vertical strip at a different density. The white lines denote successive proton cyclotron harmonics, while the black lines denote the value of the lower hybrid frequency (in units of ω_{cp}) for each value of density.	64
4.10	Simulated frequency vs magnetic field plots for the non-linear stage of the MCI. Frequency is in units of the proton cyclotron frequency. Each panel shows the variation of the \log_{10} of the spectral power in the B_z field component with frequency and magnetic field at constant density. The black and white lines indicate the position of the lower hybrid frequency (in units of the proton cyclotron frequency) and successive proton cyclotron harmonics respectively.	66
4.11	The variation of the \log_{10} of the average spectral power in the B_z field component of a given harmonic as a function of magnetic field and electron particle density. Each panel corresponds to a different simulation harmonic. Black dots denote the values of B_0 and n_e for which the lower hybrid frequency is within $\pm l/2$ of the harmonic in question, where l is the harmonic number.	67

- 4.12 Temporal evolution of ICE amplitude during an ELM crash in KSTAR plasma 11513. Time $t = 0$ refers to the centre of a $200 \mu s$ segment of radio-frequency data. The horizontal dashed lines in the spectrogram indicate energetic proton cyclotron harmonics f_{cp} at the low field side plasma edge. In addition to the main chirping feature $\lesssim 500$ MHz $\approx 20 f_{cp}$ discussed in section 4.1, we also observe a second, faint (“ghost”), feature at frequencies above the lower hybrid frequency $f_{LH} \approx 529\text{MHz} \approx 21 f_{cp}$. This additional, spectral feature is delayed in time by approximately $1\mu s$ with respect to the main chirping feature. 68
- 4.13 Left: Bicoherence of the observed fast RF signal displayed in Fig. 4.12, plotted as a function of frequency in MHz. The colour scale indicates intrinsic nonlinear coupling between waves with frequencies f_1 and f_2 , which takes values between 0 and 1. There is significant coupling in three distinct regions, discussed as (i) to (iii) in this section. Right: Bispectrum of the same RF signal. Bicoherence measures the intrinsic strength of nonlinear wave coupling, while the bispectrum measures actual nonlinear transfer of energy. Colour is plotted on a log scale. 70
- 4.14 Upper panels: Normalised bicoherence (Eq. (2.19)) of the oscillatory part of the B_z field component in three PIC simulations. This is plotted as a function of wavenumber normalised to f_{cp}/V_a where V_a is the Alfvén speed. From left to right the number density n_e in the simulations is $2.4 \times 10^{19}\text{m}^{-3}$, $1.3 \times 10^{19}\text{m}^{-3}$, and $0.8 \times 10^{19}\text{m}^{-3}$. In all three plots the most significant coupling is observed between neighbouring k values near the $k_1 = k_2$ boundary. Lower panels: Corresponding spatio-temporal Fourier transform for the three simulations. The y-axis is plotted in units of f_{cp} while the x-axis is plotted in units of f_{cp}/V_a . The horizontal black line denotes the lower hybrid frequency f_{LH} for each simulation. In addition to the expected cold plasma waves below f_{LH} , spectrally intense regions above f_{LH} can be seen in the range of frequencies corresponding to the observed “ghost” feature in Fig. 4.12. 71
- 4.15 Spatio-temporal Fourier transform of the E_x field component of the simulation which makes up the left panels of Fig. 4.14, in which $n_e = 2.4 \times 10^{19}\text{m}^{-3}$. The y-axis is plotted in units of f_{cp} while the x-axis is plotted in units of f_{cp}/V_a . The horizontal black line denotes the lower hybrid frequency f_{LH} for each simulation. 72

4.16	Left panels: reproduction of Fig. 4.4. Right, top panel: expanded view of the upper region of Fig. 4.12 in which $f \geq f_{LH}$. Right, lower panels: chirping is apparent in frequency versus number density plots for the nonlinear stage of MCI simulations. Shading indicates the \log_{10} of the spectral power in the fluctuating part of the B_z field component of each simulation.	75
4.17	Upper panel: 200MHz RF (black) and D_α (red) signals around the ELM crash during KSTAR plasma 16176. Lower panel: Temporal evolution of ICE amplitude during an ELM crash in KSTAR plasma 16176. Time $t = 0$ denotes the time at which the first derivative of the RF signal is almost discontinuous, which coincides with the peak in the RF signal, the bursting phase [Kim et al., 2018]. During windows A and B, the ICE signal shows spectral peaks at successive harmonics of the deuterium cyclotron frequency. The ICE signal during windows C and D shows more complex burst phenomena, which are discussed in Refs. [Thatipamula et al., 2016; Kim et al., 2018; Chapman et al., 2017, 2018].	76
4.18	Contour plot displaying the analytical linear MCI growth rate of the fastest growing mode as a function of pitch angle and particle energy, using a \log_{10} colour scale. Motivated by experimental observations, these calculations are restricted to frequencies below the 30th deuteron cyclotron harmonic. The linear MCI growth rate is exponentially strongest for pitch angles in the range $78^\circ < \phi < 85^\circ$. For a given pitch angle, the linear MCI is strongest at higher NBI deuteron energy; the strength of this dependence increases with pitch angle. The range of pitch angles displayed reflects the range for which the CUEBIT test particle code predicts orbits which are within the plasma boundary and traverse the ICE emitting region, see also Fig. 4.19.	79
4.19	Poloidal projection of 100keV NBI deuteron orbits calculated using the CUEBIT test particle code using a Solov'ev approximation to a typical KSTAR equilibrium. Particle orbits are initialised near the core for four different pitch angles within the range $80^\circ \leq \phi \leq 84^\circ$. Plots of full orbits are labelled a)-d) with corresponding insets to the right of each panel. Panels a), b), c), and d) show orbits with initial pitch angles of 80° , 81° , 83° , and 84° respectively.	80

4.20 Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time, from a PIC simulation with $\xi = 10^{-3}$. The traces, ordered from top to bottom at their peak (and in colour online) are: Top (red) the change in kinetic energy density of the thermal bulk plasma deuterons; second (black) change in energy density of the electrons; third (blue) the energy density of the electrostatic field E_x ; fourth (green) the energy density of the magnetic field perturbation ΔB_z ; fifth (cyan) the change in kinetic energy density of the minority energetic NBI deuterons. Time is normalised to the deuteron gyro period. The primary energy flow from the NBI deuterons is to the thermal deuterons, whose kinetic oscillation helps support the field oscillations excited by the MCI. These field oscillations include, with comparable magnitude, an electromagnetic component $(\Delta B_z)^2$ and an electrostatic component E_x^2 . The electrostatic component involves electron kinetics which are fully captured in our PIC model. The MCI saturates within five deuteron gyroperiods. The initial energy densities of the E_x field component and the fluctuating ΔB_z field component are both zero; while the initial energy densities of the electrons, thermal deuterons, and NBI deuterons are $\sim 6 \times 10^3 \text{Jm}^{-3}$, $\sim 6 \times 10^3 \text{Jm}^{-3}$, and $\sim 360 \text{Jm}^{-3}$ respectively. 83

- 4.21 Distribution of energy in the fluctuating z-component of the magnetic field ΔB_z across frequency-wavenumber space from a PIC simulation with $\xi = 10^{-2}$, $T_e = T_D = 1\text{keV}$, $B_{0z} = 1.44\text{T}$, $n_e = 2.5 \times 10^{19}\text{m}^{-3}$, and a 100keV minority NBI deuteron population. This plot is a spatio-temporal Fourier transform of the B_z field over the intervals spanning $0 \leq x \leq 50,000\lambda_D$ and $0 \leq t \leq 5\tau_{cD}$. Shading indicates the \log_{10} of the spectral density of the oscillatory part ΔB_z of the B_z field component in frequency-wavenumber space. The sweep of the fast Alfvén wave from bottom left to top right is intersected by cyclotron harmonic waves at successive deuteron harmonics. The phase velocity of the fast Alfvén wave $\simeq v_A$, and this exceeds the speed v_{NBI} of the NBI deuterons which is plotted as a blue diagonal line. Wave excitation is strongest in the wedge between v_A and v_{NBI} , and in particular where cyclotron harmonic waves intersect the fast Alfvén wave. Simulated ICE frequency spectra, such as the lower panels of Fig. 4.22, are obtained by integrating plots such as Fig. 4.21 over wavenumber. 84
- 4.22 Panels (a) and (b): measured spectral intensity of the ICE signal from KSTAR deuterium plasmas with 100keV deuteron NBI heating: (a) [Kim et al., 2018] plasma 16176 with $B_0 = 1.84\text{T}$, (b) [Thatipamula et al., 2016] plasma 11474 with $B_0 = 2.27\text{T}$. In both panels the FFT was performed using data from $t = -150\mu\text{s}$ to $t = -100\mu\text{s}$ relative to t_0 (see, for example, Fig. 4.17). Panels (c) and (d): Blue traces are outputs of PIC simulations of the spectral intensity of the fluctuating B_z field energy density, resulting from relaxation of a minority 100keV deuteron ring-beam population in thermal deuterium plasma. The simulation parameters of (c) and (d) map to (a) and (b) respectively, corresponding to the local ICE-emitting plasma in the two KSTAR experiments: $n_e = 2.5 \times 10^{19}\text{m}^{-3}$, $T_e = T_D = 1\text{keV}$. Green traces provide a noise baseline for the blue traces. They are obtained from the thermal plasma without a ring-beam, so that any spectral structure arises from the fluctuation-dissipation theorem and identifies normal modes. 85

4.23	The return loss (S_{11}) as a function of frequency of the Spiral (blue) and Bowtie (red) antennas used for the detection of signals in panels (a) and (b) respectively. The y-axis denotes the return loss and is on a log scale. For frequencies less than 150MHz, the return loss of the Bowtie antenna is close to 0dB, offering a likely explanation of the mismatch between experimental and simulation spectral peaks at low harmonic number ($l \leq 11$) seen in panels (b) and (d) of Fig. 4.22.	86
5.1	Left: time evolution of the change in energy density of particles and electric and magnetic fields as a function of time for a PIC simulation of the MCI with JET like plasma parameters and a wave propagation angle of $\theta = 90^\circ$. Right: as left but using $\theta = 89^\circ$. In both cases the parallel component of the ^3He ring beam velocity $v_{\parallel 0}$ is set to 0 and time is normalised to the ^3He gyro-period.	91
5.2	Distribution of energy in the B_z field component across frequency-wavenumber space. Shading indicates the \log_{10} of the spectral density. The temporal Fourier transform is calculated between $0\tau_{g\text{He}^3} \leq t \leq 9\tau_{g\text{He}^3}$, while the spatial Fourier transform is calculated across the entire simulation domain. The vertical axis is normalised to the ^3He cyclotron frequency while the horizontal axis is normalised to the ^3He cyclotron frequency divided by the Alfvén speed. Left: $\theta = 90^\circ$. Right: $\theta = 89^\circ$	92
5.3	Spectral intensity of the B_z field component of the MCI simulation using a temporal Fourier transform window spanning data from $0\tau_{gp} \leq t \leq 9\tau_{gp}$. The vertical axis is on a \log_{10} scale while the horizontal axis is normalised to the ^3He cyclotron frequency. Left: $\theta = 90^\circ$. Right: $\theta = 89^\circ$	92
5.4	Left: time evolution of the change in energy density of particles and electric and magnetic fields as a function of time for a PIC simulation of the MCI with JET like plasma parameters and $\theta = 89^\circ$. The parallel component of the ring beam velocity $v_{\parallel 0} = 0$ and time is normalised to the ^3He gyro-period. Right: as left except $v_{\parallel 0} = 5 \times 10^6 \text{ms}^{-1}$	93
5.5	Spectral intensity of the B_z field component of the MCI simulation with $\theta = 89^\circ$ and using a temporal Fourier transform window spanning data from $0\tau_{gp} \leq t \leq 9\tau_{gp}$. Left: the parallel component of the ^3He ring beam velocity $v_{\parallel 0} = 0$. Right: $v_{\parallel 0} = 5 \times 10^6 \text{ms}^{-1}$	93

5.6	Spectral intensity of the B_z field component of the MCI simulation with $\theta = 89^\circ$ and using a temporal Fourier transform window spanning data from $0\tau_{gp} \leq t \leq 9\tau_{gp}$. The vertical axis is on a \log_{10} scale while the horizontal axis is normalised to the ${}^3\text{He}$ cyclotron frequency. Left: the parallel component of the ${}^3\text{He}$ ring beam velocity $v_{\parallel 0} = 0$. Right: $v_{\parallel 0} = 5 \times 10^6 \text{ms}^{-1}$	94
5.7	Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time for four PIC simulations of the MCI with JET like plasma parameters. Panels a), b), c) and d) correspond to simulations with wave propagation angles equal to $90^\circ, 92^\circ, 94^\circ$ and 96° respectively. In each panel the parallel component of the ring beam velocity $v_{\parallel 0}$ is set to 0 and time is normalised to the ${}^3\text{He}$ gyro-period.	95
5.8	Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time. The traces, ordered from top to bottom at their peak are: Top (red) the change in kinetic energy density of the thermal bulk plasma deuterons; second (green) the energy density of the magnetic field perturbation ΔB_z ; third (blue) the energy density of the electrostatic field E_x ; fourth (magenta) the energy density of the magnetic field B_y ; fifth (cyan) the change in kinetic energy density of the minority energetic protons. Time is normalised to the proton gyro-period. The MCI saturates within ten proton gyro-periods.	98
5.9	Blue trace: spectral intensity of the B_z field component of the MCI simulation using a temporal Fourier transform window spanning data from $0\tau_{gp} \leq t \leq 10\tau_{gp}$. The green trace shows the power in a background thermal plasma without a minority energetic proton ring beam. Power is obtained by integrating across the entire spatial domain and is plotted on a \log_{10} scale.	99
5.10	Spectral intensity of the B_z field component of the MCI simulation as a function of time and wavevector plotted on a \log_{10} scale. Linearly unstable modes with $k > 8$ can be seen evolving in time. Three nonlinearly driven modes at $k \simeq 1.2, 2.5,$ and 7.1 , which correspond to $\omega \simeq 1, 2,$ and 7 can be seen to evolve from the nonlinear stage of the simulation at $t \approx 9\tau_{gp}$	100

5.11	Distribution of energy in the z-component of the magnetic field across frequency-wavenumber space. Shading indicates the \log_{10} of the spectral density. The temporal Fourier transform is calculated between $0\tau_{gp} \leq t \leq 18\tau_{gp}$, while the spatial Fourier transform is calculated across the entire simulation domain. The vertical axis is normalised to the proton cyclotron frequency while the horizontal axis is normalised to the proton cyclotron frequency divided by the Alfvén speed. Left: forward and backward propagating waves across a wide range of frequencies and wavenumbers. Right: an expanded view of the low frequency region of the dispersion plot. We see regions of high spectral density at proton cyclotron harmonics lying just to the right (left) of the forward (backward) propagating fast Alfvén wave.	101
5.12	Spectral intensity of the B_z field component of the MCI simulation using temporal Fourier transform windows of increasing duration. Panels a), b), c), and d) correspond to temporal Fourier transform windows of duration $10\tau_{gp}$, $12.5\tau_{gp}$, $15\tau_{gp}$, and $17.5\tau_{gp}$ respectively. All panels show data corresponding to the nonlinear stage of the simulation. The green trace shows the power in a background thermal plasma without a minority energetic proton ring beam, and hence without the excitation of the MCI. Power is obtained by integrating across the entire spatial domain and is plotted on a \log_{10} scale. . . .	102
5.13	The difference in Fourier power of the B_z field component between the $\omega = 1\omega_{cp}$ mode of an MCI simulation and the $\omega = 1\omega_{cp}$ mode of a corresponding thermal plasma simulation as a function of time. . .	103
5.14	Minority proton probability density distribution functions at different times in the simulation. The y-axis of each is normalised to its own maximum probability density and the x-axis each is normalised to its own maximum momentum. Top: p_x probability density function as a function of p_x . Middle: p_y probability density function as a function of p_y . Bottom: $p_{\perp} = \sqrt{p_x^2 + p_y^2}$ probability density function as a function of p_{\perp}	105

5.15	Minority energetic proton bulk velocity as a function of time. If the distribution was perfectly symmetric, this would always be zero. Left: Bulk velocity in the x-direction. Right: Bulk velocity in the y-direction. Velocity is normalised to the electron thermal velocity $v_{th,e}$. In both cases we see a distinct change in the bulk velocity to either a net negative or net positive flow at the time of linear saturation (just before $t = 10\tau_{gp}$. This net flow in a preferred direction persists through the non-linear stage, gradually tending to 0.	106
5.16	Bicoherence for a time $t = 7\tau_{gp}$ before linear saturation. Strong interactions are not observed and there is almost no asymmetry. . .	107
5.17	Bicoherence for a time $t = 10\tau_{gp}$, shortly after linear saturation. Very strong interactions are observed in several regions, as well as a weaker, but still significant group of interactions at the forward-backward propagating wave boundary.	108
5.18	Bicoherence for a time $t = 13\tau_{gp}$, well into the nonlinearly saturated regime of the simulation. The same strong modes as in Fig. 5.17 can be seen, but the strength of the interaction between modes at the forward-backward propagating wave boundary is stronger.	108
6.1	Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time from six PIC simulations in which the minority energetic protons are initialised using ring-beam velocity distributions with varying perpendicular velocity spreads. The perpendicular spread $v_{T\perp}$ as a fraction of the initial perpendicular velocity $v_{0\perp}$ is shown at the top of each panel. The traces, ordered from top to bottom at their peak (and in colour online) are: Top (red) the change in kinetic energy density of the thermal bulk plasma deuterons; second (green) the energy density of the magnetic field perturbation ΔB_z ; third (blue) the energy density of the electrostatic field E_x ; fourth (cyan) the change in kinetic energy density of the minority energetic protons. Time is normalised to the proton gyro period. Note the different scales on the y-axis of each plot. . . .	114

6.2	Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time from six PIC simulations in which the minority energetic protons are initialised using spherical shell velocity distributions with varying velocity spreads (shell thickness). The spread v_T as a fraction of the initial velocity v_0 is shown at the top of each panel. The ordering of the traces and the normalisation is the same as in Fig. 6.1	115
6.3	Distribution of energy in the fluctuating z-component of the magnetic field ΔB_z across frequency-wavenumber space from six PIC simulations in which the minority energetic protons are initialised using ring-beam velocity distributions with varying perpendicular velocity spreads. The spread $v_{T\perp}$ as a fraction of the initial perpendicular velocity $v_{0\perp}$ is shown at the top of each panel. The plots are a spatio-temporal Fourier transform of the B_z field component and in each case span the entire spatial domain and the length of time shown in their Fig. 6.1 counterparts. Shading indicates the \log_{10} of the spectral density of the oscillatory part ΔB_z of the B_z field component. The vertical axes are normalised to the minority proton cyclotron frequency ω_{cp} , while the horizontal axes are normalised to ω_{cp} divided by the Alfvén speed V_A	117
6.4	Spectral intensity of the fluctuating B_z field energy density, from six PIC simulations in which the minority energetic protons are initialised using ring-beam velocity distributions with varying perpendicular velocity spreads. The spread $v_{T\perp}$ as a fraction of the initial perpendicular velocity $v_{0\perp}$ is shown at the top of each panel. The vertical axes are plotted on a \log_{10} scale while the horizontal axes are normalised to the minority proton cyclotron frequency ω_{cp} . The range of both axes is identical in all panels.	118
6.5	Distribution of energy in the fluctuating z-component of the magnetic field ΔB_z across frequency-wavenumber space from six PIC simulations in which the minority energetic protons are initialised using spherical shell velocity distributions with varying velocity spreads (shell thickness). The spread v_T as a fraction of the initial velocity v_0 is shown at the top of each panel. The plots are a spatio-temporal Fourier transform of the B_z field and in each case span the entire spatial domain and the length of time shown in their Fig. 6.2 counterparts. All other details are the same as Fig. 6.3.	120

6.6	Spectral intensity of the fluctuating B_z field energy density, from six PIC simulations in which the minority energetic protons are initialised using spherical shell velocity distributions with varying velocity spreads (shell thickness). The spread v_T as a fraction of the initial velocity v_0 is shown at the top of each panel. All other details are the same as Fig. 6.4.	121
6.7	Contour plot displaying the linear growth rates of the MCI as a function of velocity spread and proton cyclotron harmonic number. Left panels: Linear growth rates are calculated numerically from a kinetic dispersion solver. Right panels: Linear growth rates calculated directly from PIC simulations. Upper panels: Linear growth rates corresponding to a minority proton ring-beam distribution function with varying perpendicular velocity spread $v_{T\perp}$. Lower panels: Linear growth rates corresponding to a minority proton spherical shell distribution function with varying velocity spread v_T . The colour bar above each panel shows the magnitude of the growth rate γ normalised to the proton cyclotron frequency ω_{cp} . In all panels the eighth proton cyclotron harmonic is omitted, as it is found to have a zero or very small linear growth rate in all cases. In the upper panels, corresponding to the ring-beam distribution, the numerical and PIC growth rates differ slightly in their maximum and minimum values, but the left and right plots are qualitatively almost identical. In the lower panels, corresponding to the spherical shell distribution, the numerical and PIC growth rates have the same maximum and minimum values, but the left and right plots are not as qualitatively similar as their ring-beam counterparts.	123
6.8	Contour plot displaying the linear growth rates of the MCI as a function of velocity spread v_T and proton cyclotron harmonic number. The minority protons are represented using a spherical shell distribution and the growth rate, calculated according to Eq. 31 of [Dendy et al., 1993], is normalised to the proton cyclotron frequency ω_{cp} . This analytically calculated growth rate is maximum around the sixth proton cyclotron harmonic and decreases as the shell thickness increases.	124

6.9	Distribution of energy in the fluctuating z-component of the magnetic field ΔB_z across time-wavenumber space. Left: From a PIC simulation in which the minority energetic protons are initialised using a ring-beam velocity distribution with no perpendicular velocity spread. Right: From a PIC simulation in which the minority energetic protons are initialised using a spherical shell velocity distribution with no velocity spread. Both plots show the spatial Fourier transform of the B_z field and span the entire spatial domain. Shading indicates the \log_{10} of the spectral density of the oscillatory part ΔB_z of the B_z field component. The vertical axes are normalised to the minority proton cyclotron frequency ω_{cp} while the horizontal axes are normalised to ω_{cp} divided by the Alfvén speed V_A	126
6.10	Time evolution of the $k \approx 8.7\omega_{cp}/V_A$ and $k \approx 15.5\omega_{cp}/V_A$ field components of a PIC simulation in which the minority energetic protons are initialised using a ring-beam distribution with no perpendicular velocity spread, corresponding to the left panel of Fig. 6.9. A moving average has been applied in time in order to smooth out the field oscillations (see the green trace in Fig. 6.1). Both modes begin to gain energy close to the start of the nonlinear stage of the simulation, meaning they are likely is driven by nonlinear interactions.	127
6.11	The square of the bicoherence b^2 of the oscillatory part of the B_z field component as a function of normalised wavenumber k from six PIC simulations in which the minority energetic protons are initialised using ring-beam velocity distributions with varying perpendicular velocity spreads. The spread $v_{T\perp}$ as a fraction of the initial perpendicular velocity $v_{0\perp}$ is shown at the top of each panel. The colour scale indicates intrinsic nonlinear coupling between waves with wave numbers k_1 and k_2 , which takes values between 0 and 1. In each plot b^2 was calculated from the full 40 proton gyro periods τ_{cp} of simulation data, the width of each successive Fourier transform was $1.25\tau_{cp}$, and the overlap of each successive Fourier transform was $0.25\tau_{cp}$. The number of independent samples is thus 32, giving a significance level of $b^2 > 1/\sqrt{32} \gtrsim 0.177$	130

- 6.12 The sum of the square of the bicoherence b^2 as a function of normalised wavenumber k_3 from six PIC simulations in which the minority energetic protons are initialised using ring-beam velocity distributions with varying perpendicular velocity spreads. The spread $v_{T\perp}$ as a fraction of the initial perpendicular velocity $v_{0\perp}$ is shown at the top of each panel. The sum is calculated from corresponding data shown in Fig. 6.11, and $k_3 = k_1 + k_2$ is the child mode resulting from strong linear coupling between two parent modes k_1 and k_2 . Only combinations of k_1 and k_2 yielding a value of $b^2 \geq 0.8$ (see Fig. 6.11) are included in the sum. A large value of Σb^2 indicates that a child mode in the vicinity of k_3 arises from strong nonlinear coupling between multiple distinct parent modes k_1 and k_2 131
- 6.13 As Fig. 6.3, except only the region with frequency ω greater than the lower hybrid frequency ω_{LH} is shown. In this region we see the child modes with wavenumber $k_3 = k_1 + k_2$ and frequency $\omega_3 = \omega_1 + \omega_2$ that arise from the strong nonlinear wave-wave interactions shown in Figs. 6.11 and 6.12. 132
- 6.14 As Fig. 6.4, except only the region with frequency ω greater than the lower hybrid frequency ω_{LH} is shown. This plot is obtained by integrating over the region shown in Fig. 6.13. In each panel, not all the peaks are uniform in height, suggesting that some of the more intense spectral peaks, e.g. $\omega = 10\omega_{cp}$ in the top left panel, owe their existence to the strong nonlinear wave-wave interactions shown in Figs. 6.11 and 6.12. This mode in particular, which is driven entirely by nonlinear wave-wave interactions, has a spectral intensity of comparable magnitude to some of the linearly driven modes shown in Fig. 6.4. Thus this mode and others like it are the direct result of the minority energetic proton population relaxing under the MCI, and not merely normal modes of the system. 133

- 6.15 The square of the bicoherence b^2 of the oscillatory part of the B_z field component as a function of normalised wavenumber k from six PIC simulations in which the minority energetic protons are initialised using spherical shell velocity distributions with varying velocity spreads. The spread v_T as a fraction of the initial velocity v_0 is shown at the top of each panel. In each panel b^2 was calculated using time series spanning the range of data displayed in the corresponding panels shown in Fig. 6.2. For instance, the bicoherence shown in the top left panel is calculated using 73 proton gyroperiods of data, the bicoherence shown in the top middle is calculated using 80 proton gyroperiods of data, and so on. The minimum significance level is thus $b^2 \gtrsim 0.13$. We note that the strength of coupling in the simulation with $v_T = 0.0v_0$ shown in the top left panel is significantly less than the rest of the simulations. 135
- 6.16 The sum of the square of the bicoherence b^2 as a function of normalised wavenumber k_3 from six PIC simulations in which the minority energetic protons are initialised using spherical shell velocity distributions with varying velocity spreads. The spread v_T as a fraction of the initial perpendicular velocity v_0 is shown at the top of each panel. The sum is calculated from corresponding data shown in Fig. 6.15, and $k_3 = k_1 + k_2$ is the child mode resulting from strong linear coupling between two parent modes k_1 and k_2 . For the upper left panel, which corresponds to a simulation with $v_T = 0.0v_0$, only combinations of k_1 and k_2 yielding a value of $b^2 \geq 0.4$ (see Fig. 6.15) are included in the sum. For the rest of the panels, that is the simulations with finite velocity spread, only combinations of k_1 and k_2 yielding a value of $b^2 \geq 0.6$ are included in the sum. 136

7.1	Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time, from a PIC simulation with $\xi_{NBI} = 10^{-3}$ and $\xi_p = 10^{-3}$. The traces, ordered from top to bottom at their first peak ($t \approx 2\tau_{cD}$), are: Top (red) the change in kinetic energy density of the thermal bulk plasma deuterons; second (green) the energy density of the magnetic field perturbation ΔB_z ; third (blue) the energy density of the electrostatic field E_x ; fourth (cyan) the change in kinetic energy density of the minority energetic fusion-born protons, which are not initialised until $t = 10\tau_{cD}$; fifth (magenta) the change in kinetic energy density of the minority energetic NBI deuterons. Time is normalised to the deuteron gyro period.	143
7.2	Distribution of energy in the fluctuating z-component of the magnetic field ΔB_z across frequency-wavenumber space from a PIC simulation with $\xi_{NBI} = 10^{-3}$ and $\xi_p = 10^{-3}$. This plot is a spatio-temporal Fourier transform of the B_z field over the intervals spanning $0 \leq x \leq 50500\lambda_D$ and $0 \leq t \leq 10\tau_{cD}$ ($20\tau_{cp}$), before the fusion-born protons have been initialised. Shading indicates the \log_{10} of the spectral density of the oscillatory part ΔB_z of the B_z field component in frequency-wavenumber space. It is apparent that no visible ICE excitation is present, due to the high levels of noise (see Sec. 4.4). As such, this dispersion relation resembles that of a thermal plasma.	144
7.3	As Fig. 7.2, except the temporal part of the Fourier transform of the B_z field is over the interval spanning the full range of the simulation $0 \leq t \leq 18.5\tau_{cD}$ ($37\tau_{cp}$). Spectrally intense regions at harmonics of the proton cyclotron frequency (even harmonics of the deuteron cyclotron frequency) are present.	145
7.4	The spectral intensity of the fluctuating B_z field energy density obtained by integrating Fig. 7.3 over wavenumber. Strong spectral peaks at harmonics of the proton cyclotron frequency (even harmonics of the deuteron cyclotron frequency) are present.	146
7.5	Left: As Fig. 7.1 but from a simulation with $\xi_{NBI} = 10^{-2}$ and $\xi_p = 10^{-2}$. Right: An inset of the left panel, focussing on the evolution of the z-component of the magnetic field B_z .	147

7.6	The spectral intensity of the fluctuating B_z field energy density from a simulation with $\xi_{NBI} = 10^{-2}$ and $\xi_p = 10^{-2}$. Power is obtained by performing a spatio-temporal Fourier transform of the B_z field over the intervals spanning $0 \leq x \leq 50500\lambda_D$ and $0 \leq t \leq 10\tau_{cD}$ ($20\tau_{cp}$) and then integrating over wavenumber. The temporal range of simulation data corresponds to times before the initialisation of the minority fusion-born protons.	148
7.7	As the Left panel of Fig. 7.6, except the temporal part of the Fourier transform of the B_z field component is over the interval spanning the full range of the simulation, $0 \leq t \leq 18\tau_{cD}$ ($36\tau_{cp}$), after the fusion-born protons have been initialised. ICE at harmonics of the proton cyclotron frequency (even harmonics of the deuteron cyclotron frequency) is more intense than ICE at harmonics of the deuteron cyclotron frequency.	149
7.8	Experimentally-measured fast RF burst spectrogram from KSTAR plasma 11513 with $B_0 = 1.99\text{T}$ and average electron number density before the ELM crash $\langle n_e \rangle = 2.6 \times 10^{19}\text{m}^{-3}$. The white dashed lines denote successive proton cyclotron harmonics. Downward step-wise frequency chirping with $\omega_{cp}/2\pi = f_{cp} \sim 25\text{MHz}$ is apparent, and some of the features are accompanied by side-band features of lesser intensity.	150
7.9	Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time in multiple PIC simulations with initial NBI deuteron energies 80keV, 140keV, and 200keV (rows), and initial helium ash temperatures 0.1MeV, 0.5MeV, and 1.0MeV (columns). The traces, ordered from top to bottom in the upper left panel are: Top (magenta) the change in kinetic energy density of the minority helium ash; second (red) the change in kinetic energy density of the thermal bulk plasma deuterons; third (green) the energy density of the magnetic field perturbation ΔB_z ; fourth (blue) the energy density of the electrostatic field E_x ; fifth (cyan) the change in kinetic energy density of the minority energetic NBI deuterons. Time is normalised to the deuteron gyro period.	151

7.10	Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time in a PIC simulation in which the initial NBI deuteron energy is 140keV, the initial helium ash temperature is 0.1MeV, $\xi_{NBI} = 10^{-3}$, and $\xi_{\alpha} = 10^{-6}$. Time is normalised to the deuteron gyro period. No appreciable energy transfer takes place between the NBI deuterons and the helium ash.	152
7.11	Perpendicular velocity distribution function of the helium ash population in multiple PIC simulations with initial NBI deuteron energies 80keV, 140keV, and 200keV (rows), and initial helium ash temperatures 0.1MeV, 0.5MeV, and 1.0MeV (columns). The y-axes are normalised to 10^{-7} , while the x-axes are normalised to the ensemble average of the magnitude of the initial velocity of the NBI deuteron population. In each panel, the distribution function was calculated using particle data outputted directly from the PIC simulation at the time of linear saturation, which corresponds to the troughs of the cyan traces shown in Fig. 7.9, and is displayed in the top right inset of each panel.	154
B.1	Top: Experimental spectrogram showing the power in the 17th proton cyclotron harmonic as a function of frequency and time. Bottom: The result of several independent simulations of the MCI which shows the power in the 17th harmonic as a function of frequency and n_e	169
B.2	Top: Experimental spectrogram showing the power in the 17th harmonic as a function of a narrow range of frequency and time. Bottom: The result of several independent simulations of the MCI which shows the power in the 17th harmonic as a function of a narrow range of frequency and n_e	170
B.3	Top: Experimental spectrogram showing the power around the 17th harmonic as a function of a narrow range of frequency and time. The black boundary to the left of the letter indicates the associated time point, the black boundaries to the left of t_i and t_f indicate the first and last time points respectively. Bottom: Power in the simulations around the 17th harmonic as a function of a narrow range of frequency and n_e . Δt indicates the width between successive time points, the letters (a) to (j) denote a mapping between n_e and t	171

B.4	<p>Top: Experimental spectrogram showing the power in the 12th harmonic as a function of frequency and time. Bottom: Power around the corresponding 12th harmonic as a function of frequency and n_e. t_i and t_f denote the time points corresponding to $n_e = 0.5 \times 10^{19} \text{m}^{-3}$ and $n_e = 0.4 \times 10^{19} \text{m}^{-3}$ respectively. Δt_i and Δt_f indicate the estimated errors in t_i and t_f respectively. The labels (a) and (b) denote a mapping between n_e and t.</p>	173
B.5	<p>Experimental spectrogram showing the power in the 11th harmonic as a function of frequency and time. t denotes the time point which in this case corresponds to $n_e = 0.5 \times 10^{19} \text{m}^{-3}$. The approximate width of the bright region of the harmonic is denoted by $2\Delta t$, where Δt is the estimated error.</p>	174

Acknowledgments

First and foremost I would like to thank my supervisors Richard Dendy and Sandra Chapman, not only for their invaluable scientific advice, but for their constant mentorship and patience.

I would also like to thank colleagues at CCFE and POSTECH, particularly Ken McClements and Gunsu Yun, both of whom helped make this thesis possible.

I am grateful to all my colleagues in the CFSA, especially Bernard and Sam for support during the early years of my PhD, and throughout.

A special thank you to the JSPS organisation and all those at Kyushu University for hosting me for three months in summer 2018, in particular Kosuga-san and Inagaki-sensei. The skills and friendships I developed in Japan will be with me for a long time.

Lastly, thank you to: Josh, Luke, Sam, and Παφαελα.

Declarations

I declare that the work presented in this thesis is my own original work, unless otherwise indicated, and has not been submitted, partially or entirely, for the attainment of a degree in this or other academic institution. All the work presented here was carried out during the period from October 2015 to March 2019 under the supervision of Prof. Richard O. Dendy and Prof. Sandra C. Chapman. Some parts of this work have been published or await to be published as indicated below:

Chapter 4

B. Chapman, R. O. Dendy, K. G. McClements, S. C. Chapman, G. S. Yun, S. G. Thatipamula, and M. H. Kim, Sub-microsecond temporal evolution of edge density during edge localized modes in KSTAR tokamak plasmas inferred from ion cyclotron emission, *Nuclear Fusion*, **57**, 124004 (2017)

B. Chapman, R. O. Dendy, S. C. Chapman, K. G. McClements, G. S. Yun, S. G. Thatipamula, and M. H. Kim, Nonlinear wave interactions generate high-harmonic cyclotron emission from fusion-born protons during a KSTAR ELM crash, *Nuclear Fusion*, **58**, 096027 (2018)

B. Chapman, R. O. Dendy, S. C. Chapman, K. G. McClements, G. S. Yun, S. G. Thatipamula, and M. H. Kim, Interpretation of suprathreshold emission at deuteron cyclotron harmonics from deuterium plasmas heated by neutral beam injection in the KSTAR tokamak, *Submitted to Nuclear Fusion*

Chapter 5

K.G. McClements, A. Brisset, **B. Chapman**, S. C. Chapman, R. O. Dendy, P. Jacquet, V. G. Kiptily, M. Mantsinen, and B. C. G. Reman, Observations and modelling of ion cyclotron emission observed in JET plasmas using a sub-harmonic arc detection system during ion cyclotron resonance heating, *Nuclear Fusion*, **58**, 096020 (2018)

Chapter 6

B. Chapman, R.O. Dendy, S.C. Chapman, L Holland and S W A Irvine, Theory and simulation of ion cyclotron emission from energetic ion populations with a spherical shell velocity-space distribution, *In preparation for submission to Plasma Physics and Controlled Fusion*

Abstract

In this thesis I present particle in cell (PIC) simulations of ion cyclotron emission (ICE). ICE comprises suprathermal radiation in the ion cyclotron frequency range, whose spectrum peaks at successive local cyclotron harmonics of the emitting energetic ion population. ICE has previously been observed in all large toroidal MCF plasmas [Dendy and McClements, 2015; McClements et al., 2015]. ICE is caused by a collective instability, which in its linear phase corresponds to the magnetoacoustic cyclotron instability (MCI). The passive, non-invasive character of ICE measurements, suggest ICE is an attractive way forward for future energetic ion measurements in ITER. In the simulations in this thesis we use the EPOCH particle-in-cell code to solve the self-consistent Maxwell-Lorentz system of equations for fully kinetic electrons and thermal background ions, together with the minority energetic ion distribution that drives the primary ICE.

We first perform a detailed quantitative comparison between fusion born proton driven chirping ICE observed during KSTAR ELM crashes and fully nonlinear direct numerical simulations of the MCI. We find good quantitative agreement between the simulated and observed spectra, to the extent that the simulations can be used to infer fast ($\sim \mu s$) time scale dynamics of the local electron number density in the emitting region. We then extend this study to determine the origin of a faint, time delayed proton chirping feature observed in one of the KSTAR plasma pulses. We do this using bicoherence analysis of both experimental and simulation data. We then run MCI PIC simulations of the pre ELM crash “steady state” ICE observed on KSTAR, which is believed to be driven by neutral beam injected (NBI) deuterons. PIC simulations of MCI excited ICE in JET and ASDEX Upgrade (AUG) plasmas are then discussed, and we show that AUG observations of the fundamental ICE harmonic can only be explained in terms of the MCI if nonlinear wave-wave interactions between higher harmonics are taken into account. Motivated by recent observations of ICE in the core region of several tokamaks, including AUG and DIII-D, we then compare MCI simulations using two types of energetic ion distribution function, a spherical shell of varying thickness, and a ring beam of varying width. It is found that both distribution functions lead to MCI excited waves, and their nonlinear properties are discussed.

Chapter 1

Introduction

1.1 Nuclear Fusion

The most ubiquitous example of nuclear fusion is stellar nucleosynthesis, whereby main sequence stars are created due to the fusion of light elements into heavier ones. The Sun is our closest natural thermonuclear reactor, providing Earth with a seemingly limitless supply of energy. The aim of fusion energy research is to be able to reproduce these reactions and harness this energy on Earth.

In the past few decades, multiple approaches have been employed in an attempt to obtain terrestrial fusion power. Four common methods of achieving this which still benefit from contemporary research are:

1. Magnetic Pinches,
2. Inertial Electrostatic confinement,
3. Inertial confinement fusion,
4. Magnetic confinement fusion.

The focus of this thesis is magnetic confinement fusion, I will therefore only briefly discuss methods one to three.

In a magnetic pinch, a strong current is sent in a particular direction through a gas which is then ionised. The current forms a strong magnetic field which “pinches” the plasma forming filamentary current strands. Eventually, these filaments bunch together, forming dense, magnetically confined hot spots which result in fusion. In inertial electrostatic confinement, a voltage difference focuses charged particles radially in either a spherical or cylindrical geometry. Ions accelerate down

the electrostatic potential well and converge at the origin, creating a high-density fusion core.

Inertial confinement fusion (ICF) is one of two main branches of fusion research. Energy from a laser is deposited on the outer layer of a target, which is typically a pellet containing a mix of deuterium and tritium. The heated outer layer ionises and ablates, the reaction drives a pressure shock wave into the centre of the pellet. If a set of shock waves are sufficiently powerful, the fuel at the center can be compressed and heated so much that fusion reactions occur.

Magnetic confinement fusion (MCF) is the other main branch of fusion research. It relies on powerful magnetic fields to confine the plasma within a chamber. Since the 1950s, several devices have been created in the attempt to achieve sufficient magnetic confinement. Much of the early research was focused around the magnetic mirror device. Confinement of plasma within this device relies on the invariance of the magnetic moment μ

$$\mu = \frac{1}{2} \frac{mv_{\perp}^2}{B}, \quad (1.1)$$

where m is the mass of the charged particle, v_{\perp} is the perpendicular component of its velocity, and B is the magnitude of the magnetic field strength. As a particle moves from a region of weak magnetic field to strong magnetic field, it experiences an increasing B . To conserve μ , the perpendicular component of the velocity must increase. Since the total energy of the particle must remain constant, the parallel component of the velocity v_{\parallel} must decrease. If B is high enough at the far end of the mirror, then $v_{\parallel} \rightarrow 0$ and the particle is reflected back [Chen, 1984]. This concept is illustrated in Fig. 1.1.

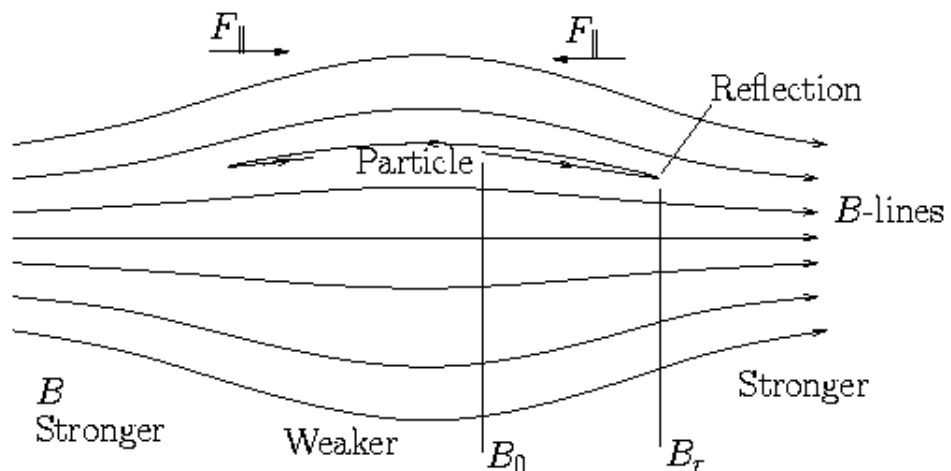


Figure 1.1: A plasma trapped between magnetic mirrors [Hutchinson, 2001].

There are problems with confining a plasma using this method, the result is that the confinement time is not sufficiently long enough for fusion energy production. However, some of the basic principles of this configuration are still used in magnetic confinement devices today. This concept is not directly relevant to this thesis, but is here to provide the reader with a simple picture of a magnetic confinement device.

Other magnetic confinement devices include the reverse field pinch, and the stellarator; but the most well studied, and arguably the most successful to date, is the Tokamak. This is a toroidal “doughnut” shaped chamber in which helical magnetic field lines, their vector pointing predominantly in the toroidal direction, confine the plasma in a stable equilibrium. This concept will be discussed in more detail in later sections.

The choice of reactants in a thermonuclear fusion device is an important issue. One requires the fusion reaction to produce a significant amount of energy, but also that the probability of the reaction occurring is sufficiently high, for a range of temperatures that can be achieved in the laboratory. The fusion of deuterium with tritium has been identified as the most efficient fusion reaction in magnetic confinement devices as it fulfils the above requirements. The reaction is



The neutron carries away 14.1MeV of the energy, while the alpha particle carries 3.5MeV of energy. Not only is this reaction promising in terms of its energy yield, it is also economically viable. Deuterium is not radioactive, occurs naturally in nature, and is a virtually inexhaustible resource. Tritium does not occur naturally in nature, as it has a half-life of approximately 12.3 years. However, the neutrons in the above reaction are electrically neutral so are not confined by magnetic fields. They escape the plasma and are absorbed by a blanket surrounding the container walls. In future tokamak devices such as the International Thermonuclear Experimental Reactor (ITER) [Hawryluk et al., 2009], these blankets will contain lithium, a relatively abundant resource, which can capture the neutron and breed tritium.

The remainder of this chapter is designed to give the reader an overview of the main elements of plasma physics used in the work presented in this thesis. Fundamentals of plasma physics including particle motion and plasma waves are introduced in Sec. 1.2, followed by a brief account of Tokamak plasmas including instabilities in Sec. 1.3. In Sec. 1.4 I will discuss Ion cyclotron emission (ICE) and the magnetoacoustic cyclotron instability (MCI) which form the basis of the work in this thesis.

1.2 Plasma physics

1.2.1 Single particle effects

Plasma dynamics can be described in various ways depending on the temporal and spatial scales at which the plasma phenomena under study occur. The theory of ideal magnetohydrodynamics (MHD) is used when one wishes to treat the plasma as a magnetised fluid.

Ideal MHD is a single fluid model of the plasma such that $T_e = T_i$. The plasma is assumed to be quasi-neutral and sufficiently collisional so that the particle distribution function can be represented by a Maxwellian. Such collisions frequently occur in magnetically confined plasma and the mean free path is often large compared to the system length. This treatment is valid only on spatial scales L longer than the single-particle characteristic scales, that is the Debye length λ_D (the typical distance over which the electric field in a plasma is shielded), and the electron and ion Larmor radii $r_{L,e}$ and $r_{L,i}$. Similarly, the temporal scales of the problem must be longer than the inverse plasma frequency ω_p^{-1} and the electron and ion gyro-periods $\tau_{e/i} = \Omega_{e/i}^{-1}$ (see Eq. 1.4 and Eq. 1.5. Note that assuming $L \gg \lambda_D$ amounts to assuming quasi-neutrality. If one would like to account for the dynamics of individually charged particles then the approximations of ideal MHD do not apply. In this instance it is more appropriate to use a kinetic description to model the plasma. The work presented in this report is carried out using the latter description, in which a particle with charge q experiences the Lorentz force

$$F = m \frac{d\mathbf{v}}{dt} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}), \quad (1.3)$$

where m refers to the mass of the particle, and \mathbf{v} the velocity of the particle. In the absence of an electric field \mathbf{E} , and with a magnetic field \mathbf{B} that is uniform along the z -axis, it is straightforward to decompose Eq. 1.3 into its spatial components and obtain a description of a particle's velocity and position. Upon doing so, it is the case that the motion of the particle consists of circular motion perpendicular to the magnetic field, and a uniform velocity along it. Therefore, a particle follows a helical orbit, gyrating around a guiding centre which follows the magnetic field lines for a homogeneous, straight magnetic field and no electric field. The frequency of the perpendicular motion is known as the cyclotron/gyro frequency

$$\Omega_c = \frac{qB}{m}, \quad (1.4)$$

this type of orbit is known as a Larmor orbit and has a radius (from the guiding

centre) known as the Larmor radius

$$r_L = \frac{v_\perp}{\Omega_c} = \frac{mv_\perp}{qB}. \quad (1.5)$$

In a tokamak, the magnetic field configuration is not straight, nor is the plasma homogeneous. This, among other things, gives rise to a plethora of particle drifts, which are briefly discussed in Sec. 1.2.

In the kinetic theory of plasmas, the electromagnetic fields are evolved self-consistently using the complete set of Maxwell's equations

$$\nabla \cdot \mathbf{E} = \frac{\sigma}{\epsilon_0}, \quad (1.6)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (1.7)$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad (1.8)$$

$$\nabla \times \mathbf{B} = \mu_0 \left(\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right), \quad (1.9)$$

where σ is the total charge density, and \mathbf{J} is the total current density. It is often more convenient to adopt a statistical approach to describe the plasma. To do so, one requires a velocity distribution for each species. Such a distribution has several independent variables and is given by: $f(\mathbf{r}, \mathbf{v}, t) = f(x, y, z, v_x, v_y, v_z, t)$. This gives rise to the so-called Vlasov equation

$$\frac{\partial f_j}{\partial t} + \mathbf{v}_j \cdot \nabla_r f_j + \frac{q_j (\mathbf{E} + \mathbf{v}_j \times \mathbf{B})}{m_j} \cdot \nabla_v f_j = 0. \quad (1.10)$$

Equations 1.6 through to 1.10 are collectively known as the Vlasov-Maxwell system of equations. By taking the moments of these equations for each particle species, one can derive the MHD equations for each plasma component. If we consider a plasma consisting only of two plasma species, electrons and one ion species, we can combine the multispecies fluid equations to derive the *ideal* MHD equation. The chief simplification, among others, is that $m_e \ll m_i$, such that terms of order $\sim m_e/m_i$ can be neglected.

1.2.2 Plasma waves

With the aid of the dielectric tensor ϵ , Maxwell's equations can be solved for plane wave solutions [Stix, 1992; Dendy, 1990]. We first define the convenient dimensionless vector $\mathbf{n} = \frac{\mathbf{k}c}{\omega}$, and θ to be the angle between z-axis and the direction of wave propagation, which is assumed to lie in the x-y plane. Under these conditions, the plasma wave solutions for a cold, collision-less plasma, with density and magnetic field static in time and homogeneous in space, are given by the following determinant

$$\begin{vmatrix} S - n^2 \cos^2(\theta) & -iD & n^2 \cos(\theta) \sin(\theta) \\ iD & S - n^2 & 0 \\ n^2 \cos(\theta) \sin(\theta) & 0 & P - n^2 \sin^2(\theta) \end{vmatrix} = 0. \quad (1.11)$$

Denoting the cyclotron frequency for a particle of species "s" by Ω_{cs} , and the plasma frequency for a particle belonging to the same species by ω_{ps} , we have

$$\begin{aligned} D &= \frac{1}{2}(R - L), \\ S &= \frac{1}{2}(R + L), \\ P &= 1 - \sum_s \frac{\omega_{ps}^2}{\omega^2}, \\ R &= 1 - \sum_s \frac{\omega_{ps}^2}{\omega(\omega + \Omega_{cs})}, \\ L &= 1 - \sum_s \frac{\omega_{ps}^2}{\omega(\omega - \Omega_{cs})}. \end{aligned} \quad (1.12)$$

It can be shown that the solutions are given by

$$n^2 = \frac{B \pm F}{2A}, \quad (1.13)$$

where

$$\begin{aligned} B &= RL \sin^2(\theta) + PS(1 + \cos^2(\theta)) \\ F^2 &= (RL - PS)^2 \sin^4(\theta) + 4P^2 D^2 \cos^2(\theta) \\ A &= S \sin^2(\theta) + P \cos^2(\theta). \end{aligned} \quad (1.14)$$

Using the above, one is able to obtain a dispersion relation between ω and \mathbf{k} , which

is not necessarily single valued. The dispersion relation of a deuterium plasma with $n_e = 2.5 \times 10^{19} \text{m}^{-3}$ and $B = 2.5 \text{T}$, typical of a tokamak plasma, is plotted in Fig. 1.2 for frequencies in the ion cyclotron range only. The red trace shows waves propagating parallel to the magnetic field ($\theta = 0$), while the blue trace shows waves propagating perpendicular to the magnetic field ($\theta = 90$).




Figure 1.2: Cold plasma dispersion relation for parallel propagating (red) and perpendicularly propagating (blue) waves. Only waves in the ion cyclotron range of frequencies are shown. The y-axis is normalised to the deuteron cyclotron frequency, while the x-axis is normalised to the deuteron cyclotron frequency divided by the Alfvén speed. The horizontal dashed line denotes the deuteron cyclotron frequency, the diagonal dashed line denotes the Alfvén speed, and the horizontal solid line denotes the lower hybrid frequency for perpendicular propagation.

In the case of perpendicular propagation, there is a resonance frequency shown by the horizontal solid black line. For this set of plasma parameters, this frequency is at $\omega \sim 33\Omega_{cD}$ where Ω_{cD} is the deuteron cyclotron frequency. This is known as the lower hybrid frequency and for strictly perpendicular propagation in the cold

plasma limit, it is given by

$$\omega_{LH} \approx \left(\omega_{pi}^{-2} + (\Omega_{ce}\Omega_{ci})^{-1} \right)^{-1/2}, \quad (1.15)$$

when $\omega_{pi} \gg \Omega_{ci}$. Here ω_{pi} is the ion plasma frequency and Ω_{ci} and Ω_{ce} are the ion and electron cyclotron frequencies respectively. The frequency range between ω_{LH} and the higher frequency

$$\omega_2 = -\frac{\Omega_{ce}}{2} + \sqrt{\frac{\Omega_{ce}^2}{4} + \omega_{pe}^2}, \quad (1.16)$$

defines a region of evanescence, that is, waves with frequency $\omega_{LH} < \omega < \omega_2$ cannot propagate. Note that ω_2 is in the electron cyclotron frequency range and hence not shown in Fig. 1.2. See Fig. 4.4 of Ref. [Cairns, 1985] for a helpful diagram. For propagation angles such that $\cos^2(\theta) \lesssim \frac{m_e}{m_i}$, the corresponding ‘‘lower-hybrid’’ frequency for waves propagating at oblique angles is given by [Verdon et al., 2009]

$$\omega = \omega_{LH} \left(1 + \frac{m_i}{m_e} \cos^2(\theta) \right)^{1/2}. \quad (1.17)$$

The above statement of evanescence only applies strictly to electrostatic, cold, linear, perpendicularly propagating waves. We shall see later that there are cases in which waves can exist in this region of frequency space.

Table 1.1 contains information about the cold plasma waves whose dispersion relation can be obtained by solving Eq. 1.11. It is useful to note that electrostatic waves are purely longitudinal, that is $\mathbf{k} \parallel \mathbf{E}$, and therefore if an electrostatic wave is propagating along the x-direction, its properties can be deduced by examining the E_x field component. Electromagnetic waves are predominantly transverse (sometimes with a longitudinal component), with $\mathbf{k} \perp \mathbf{E}$, and consequently, they manifest themselves in the E_y , E_z , B_y , and B_z field components (in both cases the perturbation to the plasma motion is in the direction of \mathbf{E}). Here, \mathbf{B}_0 refers to the uniform background magnetic field, not the total magnetic field including the perturbation. All electromagnetic waves shown in table 1.1 are transverse, with the exception of the X-wave which has a longitudinal component.

¹The dispersion relation of this wave is modified by the inclusion of warm plasma effects.

²The dispersion relation is identical to that of a light wave in a magnetised plasma.

³The dispersion relation for this wave splits into two branches.

⁴The dispersion relation for this wave splits into two branches, the lower of which is commonly referred to as the ‘‘whistler’’ branch/wave.

Table 1.1: Waves in cold uniform plasmas.

Name	Particle species	Character and Conditions
Plasma oscillation ¹	Electrons	Electrostatic, $\mathbf{B}_0 = 0$ or $\mathbf{k} \parallel \mathbf{B}_0$
Upper hybrid oscillation	Electrons	Electrostatic, $\mathbf{k} \perp \mathbf{B}_0$
Light wave	Electrons	Electromagnetic, $\mathbf{B}_0 = 0$
O-wave ²	Electrons	Electromagnetic (Plane polarised with $\mathbf{E} \parallel \mathbf{B}$), $\mathbf{k} \perp \mathbf{B}_0$
X-wave ³	Electrons	Electromagnetic, Transverse and Longitudinal, $\mathbf{E} \perp \mathbf{B}$, $\mathbf{k} \perp \mathbf{B}_0$
R-wave ⁴	Electrons	Electromagnetic (Right circular polarisation), $\mathbf{k} \parallel \mathbf{B}_0$
L-wave	Electrons	Electromagnetic, (Left circular polarisation), $\mathbf{k} \parallel \mathbf{B}_0$
Lower hybrid oscillation	Ions	Electrostatic, $\mathbf{k} \perp \mathbf{B}_0$
EM Ion cyclotron wave	Ions	Electromagnetic, $\mathbf{k} \parallel \mathbf{B}_0$

The terminology surrounding waves with frequencies well below the ion cyclotron frequency is often misunderstood. These waves have been omitted from table 1.1, and we discuss them in more detail here instead.

We first define $k = |\mathbf{k}|$ as the magnitude of the wavevector, and the Alfvén velocity $v_A = B/\sqrt{\mu_0\rho}$ where $B = |\mathbf{B}|$ is the magnitude of the magnetic field, and ρ is the mass density of the plasma. Sticking with the cold plasma picture, and deriving wave modes using the dielectric tensor, one can prove the existence of two other low frequency wave modes which are associated with ion dynamics. Assuming $v_A \ll c$ (which is true for most conventional tokamak plasmas), we have the slow Alfvén wave, with dispersion relation $\omega \simeq v_A k \cos(\theta) = v_A k_{\parallel}$, and the fast Alfvén wave with $\omega \simeq v_A k$. The “fast” and “slow” terminology refers to the relative phase speeds of the waves. For parallel propagation, the fast and slow Alfvén waves are indistinguishable from one another, and are commonly referred to as “shear” Alfvén waves. For perpendicular propagation, the slow wave ceases to exist, while the fast wave remains with dispersion relation $\omega = v_A k_{\perp}$ (we have $k = k_{\perp}$) and is commonly referred to as the “compressional” Alfvén wave. Thus, readers of Alfvén wave literature may often find that the words “slow” and “shear” are used interchange-

ably, as are the words “fast” and “compressional”. Physically, the perturbation of the magnetic field associated with shear Alfvén waves is directed *perpendicular* to the background magnetic field, meaning they perturb the magnetic field *direction*. In contrast, the magnetic field perturbation associated with compressional Alfvén waves is directed *parallel* to the background magnetic field, meaning they perturb the magnetic field *strength*. Incidentally, for parallel propagation, the shear Alfvén wave smoothly changes into the ion cyclotron wave discussed above.

Alternatively, one can derive the dispersion relations for these low frequency waves using the equations of MHD instead of the cold plasma dielectric tensor. The key difference is that the MHD picture includes the effects of finite plasma pressure. Mathematically, the inclusion of plasma pressure leads to terms involving the ion acoustic sound speed v_s . There are three different types of Alfvénic waves that can propagate through an MHD plasma. The first is the shear Alfvén wave, which is identical to the shear/slow Alfvén wave described in the preceding paragraph, having the same dispersion relation $\omega \simeq v_A k \cos(\theta)$. The second two waves are known as the fast magnetosonic wave and the slow magnetosonic wave and have dispersion relations $\omega = kv_+$ and $\omega = kv_-$ respectively, where

$$v_{\pm} = \left(\frac{1}{2} (v_A^2 + v_s^2) \pm \frac{1}{2} \sqrt{(v_A^2 + v_s^2) - 4v_s^2 v_s^2} \right)^{1/2}. \quad (1.18)$$

Again, the ‘fast’ and ‘slow’ terminology refers to the relative phase speeds of the waves. The word “magnetoacoustic” is sometimes used in place of “magnetosonic”, and, confusingly, the word magnetosonic/magnetoacoustic is often dropped altogether, meaning the waves are frequently referred to as the fast and slow waves. It is important to note that under these circumstances, the fast and slow waves here differ from the fast and slow waves identified in the previous paragraph, by the inclusion of finite pressure (thermal) effects. The relationship between these two instances of fast and slow waves, and the reason for the confusing terminology can be easily understood in the limit of zero plasma pressure, in which the MHD description reverts to the cold plasma dielectric description. As the pressure tends to zero, the ion acoustic speed v_s also tends to zero, and the slow magnetosonic wave ceases to exist. The dispersion relation for the fast magnetosonic wave reduces to $\omega \simeq kv_A$, which is that of the fast wave in the cold plasma description. If we substitute $\theta = 90^\circ$ into Eq. 1.18, we obtain $\omega^2 = (v_s^2 + v_A^2) k_{\perp}^2$ (we have $k = k_{\perp}$). This is the cold plasma compressional Alfvén wave modified by the presence of plasma pressure.

By now hopefully it is clear why the names fast/fast magnetoacoustic/fast magnetosonic/compressional are used almost synonymously to describe perpendicularly propagating low frequency waves, and why the names slow/slow magnetoacoustic/slow magnetosonic/shear are used almost synonymously to describe parallel propagating low frequency waves. Finally, in an attempt to clear up any remaining confusion surrounding the literature, it should be noted that the expression for the dispersion relation of the perpendicularly propagating fast magnetosonic wave, $\omega^2 = (v_s^2 + v_A^2) k^2$, assumes $v_A \ll c$. If one doesn't assume this, we arrive at a slightly different expression for the dispersion relation: $\frac{\omega^2}{k^2} = c^2 \frac{v_s^2 + v_A^2}{c^2 + v_A^2}$. It is easy to see how this reverts back to our previous expression in the limit $v_A \ll c$.

Using kinetic theory, the study of plasma waves can be taken even further by calculating the warm plasma dielectric tensor, which leads to waves such as: the ion acoustic/sound wave, electron and ion Bernstein waves, and the electrostatic ion cyclotron wave (not to be confused with its cold plasma electromagnetic counterpart). The properties of these three waves are summarised in 1.2. For this thesis, it is sufficient to say that harmonics of the ion cyclotron frequency are a warm plasma (finite Larmor radius) effect. Finally, the inclusion of inhomogeneities, geometrical effects, and nonlinear effects, can result in many more plasma waves, as well as modifying the properties of some of the waves described above.

Table 1.2: Waves in warm uniform plasmas.

Name	Particle species	Character and Conditions
ES Ion cyclotron wave	Ions	Electrostatic, $\mathbf{k} \perp \mathbf{B}_0$ (nearly)
Ion acoustic/sound wave	Ions	Electrostatic, $\mathbf{B}_0 = 0$ or $\mathbf{k} \parallel \mathbf{B}_0$
Electron/Ion Bernstein modes	electrons/ions	Electrostatic, $\mathbf{k} \perp \mathbf{B}_0$

1.3 Tokamak plasmas

This section is intended to give the reader an overview of the elements of tokamak physics required to understand and interpret this work. It starts with a discussion of the tokamak with an explanation of how the plasma is confined. Then, a brief introduction to particle orbits in a tokamak is presented, ending with a summary of “Edge Localized Modes” (ELMs), which is relevant to the work presented in chapter 4.

1.3.1 Confinement

In conventional tokamaks, such as JET, the plasma is confined by surrounding the chamber with a series of toroidal magnetic field coils. These coils carry a current I_θ which circulates in the poloidal direction, creating a toroidal component to the magnetic field $B_\phi \propto 1/R$, where R is the major radius of the tokamak. Alternatively, the plasma can be confined by passing a current-carrying rod up through the centre of the torus. The rod current creates a toroidal component of the magnetic field $B_\phi \propto 1/R$. This is known as a spherical tokamak, e.g. MAST [Chapman et al., 2015]. Due to particle drifts, a toroidal magnetic field alone is insufficient to maintain pressure in the plasma. As $B_\phi \propto 1/R$, there is a gradient $\nabla\mathbf{B}$ pointing towards the symmetry axis. This generates a “grad B” drift which is perpendicular to both \mathbf{B} and $\nabla\mathbf{B}$, which means it is directed upwards for ions, and downwards for electrons. Charge separation gives a vertical electric field \mathbf{E} , which gives rise to an “E cross B” drift, causing both ions and electrons to drift towards the vessel wall. This results in a loss of confinement, but can be remedied by introducing a poloidal field B_θ . Such drifts are not directly relevant to this work which is why they have only been mentioned in passing. A wealth of information regarding these drifts can be found in Ref. [Wesson, 2004].

With a poloidal field at the top of the tokamak cross-section, the trajectory of a charged particle in the poloidal plane follows a \mathbf{B} field line, but grad B drifts upwards away from the flux surface. At the bottom of the cross section, the grad B drift is still upwards, but this takes the particle back to the flux surface, meaning there is no net drift, no \mathbf{E} , and hence no loss of confinement.

The poloidal magnetic field is produced by a solenoid in the centre of the tokamak. Currents in the solenoid windings produce an internal vertical magnetic field. Changing the current in these windings changes the vertical magnetic field; this induces a toroidal electric field in the plasma, which leads to a change in flux and hence a toroidal current I_ϕ in the plasma. This current generates the poloidal

magnetic field B_θ as well as providing for Ohmic heating.

The poloidal magnetic field is larger on the inside of the tokamak and hence the magnetic pressure there is larger. This steers the plasma ring towards an increased major radius [Wesson, 2004]. For this reason, a third set of stabilising coils are needed. These coils wrap around the outside of the tokamak in the toroidal direction and produce a vertical magnetic field which adds to the poloidal magnetic field produced by the plasma current, contributing slightly to the overall poloidal magnetic field. If the toroidal magnetic field is sufficiently high, only a small vertical magnetic field is required to stabilise the plasma.

The use of the solenoid to provide the poloidal magnetic field highlights one drawback of the tokamak: one cannot ramp the coil/solenoid current indefinitely, so with this way of operating, the current is pulsed. The total magnetic field is given by the vector sum of all magnetic field components. This results in magnetic field lines that twist around the torus. The electromagnetic coil arrangement and the resulting fields in a conventional tokamak are shown in Fig. 1.3.




Figure 1.3: The electromagnetic coil configuration of a conventional tokamak showing the resulting magnetic field [EuroFusion].

1.3.2 Fast particles

There are three main sources of fast particles in tokamak plasmas:

1. Fusion born products.
2. Neutral beam injected ions.
3. Ion cyclotron resonance heated (ICRH) ions.

Ion cyclotron resonance heating occurs when a compressional Alfvén / fast magnetosonic wave is launched into the plasma using an external antenna. This wave has a dominantly radial (\sim perpendicular) wavenumber and must tunnel through a vacuum layer at the edge of the tokamak until it reaches a region of plasma with a sufficiently high density so it can propagate and heat the plasma [Wesson, 2004]. It is possible to choose the frequency of this wave such that it is in the ion cyclotron range of frequencies (ICRF), and therefore able to resonate and transfer energy to plasma ions [Cairns, 1996].

In neutral beam injection (NBI), energetic neutral atoms are injected into the plasma and then become ionised through collisions. The resulting ions and electrons have the same velocity and are now both confined by the tokamak magnetic field. As the ions are much more massive, almost all the initial energy carried by the neutral is now carried by the ions. These ions (and electrons) can then collide with other particles in the plasma transferring energy.

Fusion born ions are perhaps the most obvious source of fast ions in a thermonuclear fusion reactor. These ions are born in the core of the plasma where the density and temperature is at its highest. In many cases, particularly in small aspect ratio devices, these are promptly lost from the plasma, diffusing rapidly across the magnetic field lines towards the vessel wall. Ideally, these ions would impart their energy on the slower, “bulk” ions before they are lost [Cook et al., 2017].

In some cases, the fast ions described above can release their energy in the form of electromagnetic waves. This will be important in the next section, but first we take a moment to discuss the orbits of such fast ions.

The fast ions discussed above are subject to the tokamak drift effects outlined in Sec. 1.3.1. The orbits of the ions can be broadly categorised as being either trapped, in which the toroidal velocity of the ion reverses direction at a turning point [Heidbrink and Sadler, 1994; Gorelenkov et al., 2014], similar to what is shown in Fig. 1.1; or passing, in which the orbits travel in one toroidal direction only.




Figure 1.4: Poloidal cross section showing the cylindrical coordinate system used to model a tokamak. [Freidberg, 1987; Sauter and Medvedev, 2013]

Using a right handed cylindrical coordinate system (R, φ, z) shown in Fig.1.4, the equilibrium magnetic field can be defined in terms of a poloidal flux function $\psi(R, z)$

$$\mathbf{B} = \underbrace{\frac{1}{R} \frac{\partial \psi}{\partial z}}_{B_R} \mathbf{e}_R + B_\varphi(R, z) \mathbf{e}_\varphi - \underbrace{\frac{1}{R} \frac{\partial \psi}{\partial R}}_{B_z} \mathbf{e}_z. \quad (1.19)$$

The equations of motion in cylindrical coordinates are given by

$$\begin{aligned} \frac{dv_R}{dt} &= \frac{q}{m} (v_\varphi B_z - v_z B_\varphi) + \frac{v_\varphi^2}{R}, \\ \frac{dv_\varphi}{dt} &= \frac{q}{m} (v_z B_R - v_R B_z) + \frac{v_\varphi v_R}{R}, \\ \frac{dv_z}{dt} &= \frac{q}{m} (v_R B_\varphi - v_\varphi B_R), \end{aligned} \quad (1.20)$$

where q and m are the charge and mass of the particle respectively. Provided the poloidal flux function ψ is specified, the orbit of a given particle can be calculated. In general, ψ must be computed numerically using the Grad-Shavarnov equation [Wesson, 2004], which solves the static equilibrium force balance condition $\nabla P = \mathbf{J} \times \mathbf{B}$; where the current density \mathbf{J} , the magnetic field \mathbf{B} , and the pressure P all depend on the poloidal flux function. Under certain conditions, it is possible to specify ψ analytically. One such solution is that obtained by Solov'ev [Solov'ev, 1968], and is written in a convenient form proposed by Freidberg [Freidberg, 1987]

$$\psi = \psi_0 \left[\frac{\gamma}{8} \{ (R^2 - R_0^2)^2 - R_b^4 \} + \frac{1 - \gamma}{2} R^2 z^2 \right]. \quad (1.21)$$

Where R_0 is the plasma major radius, R_b is related to the minor radius, a , by $R_b = \sqrt{2R_0a}$, γ is related to the plasma elongation κ by $\gamma = \frac{\kappa^2}{1+\kappa^2}$, and ψ_0 is a constant which can be approximated by $\psi_0 \approx B_0/2q$. The central magnetic is denoted by B_0 and $q = \frac{d\Phi}{d\Psi}$ is the plasma safety factor which is defined as the rate of change of toroidal flux Φ with poloidal flux Ψ (the ratio between the number of times a magnetic field line traverses the toroidal direction to the number of times it traverses the poloidal direction). This analytical solution is typically applied to tight aspect ratio devices such as the MAST tokamak [Chapman et al., 2015; McClements and Fredrickson, 2017], however it is still useful as a means of demonstrating ion orbital behaviour in devices with larger aspect ratios, such as JET. Using the above equations, the orbit of a particle can be computed with relative ease, and is done here using the ‘‘CUEBIT’’ orbit code [Hamilton et al., 2003]. Figure 1.5 plots 3.5MeV alpha particle orbits corresponding to a JET-like equilibrium. The central magnetic field is set to 3T and the safety factor set to 2.5. In the left panel, the alpha particle was launched with initial position $(R,Z) = (3.1\text{m},0\text{m})$, and an initial velocity vector such that the component of velocity perpendicular to the local magnetic field is greater than the velocity component parallel to the magnetic field. The right panel shows the orbit of an alpha particle launched at $(R,Z) = (3.1\text{m},1.0\text{m})$, but this time, its parallel velocity component is greater in magnitude than its perpendicular component. The right panel shows a passing orbit, while the left panel shows a trapped, deeply passing, ‘‘banana’’ orbit. That is, its poloidal projection traces out a banana-like shape, and its radial excursion is such that it passes close to the tokamak edge. Such trapped particles are of great importance to this thesis, as they are responsible for ICE in the tokamak edge [Dendy and McClements, 2015; McClements et al., 2015; Gorelenkov, 2016].




Figure 1.5: Left: Trapped orbit of a 3.5MeV α particle in a JET-like equilibrium. Right: Passing orbit of a 3.5MeV α particle in a JET-like equilibrium.

1.3.3 Edge Localized Modes (ELMS)

Understanding the physics of edge localised modes (ELMs) [Zohm, 1996; Loarte et al., 2003; Kamiya et al., 2007; Leonard, 2014] in MCF plasmas is crucial for the design of future fusion power plants such as ITER [Kessel et al., 2009]. ELMs are periodic burst of ejected plasma and heat associated with high confinement, “H-mode” plasmas [Wagner et al., 1982]. ELMs are detrimental to plasma confinement because of the large heat fluxes to the vessel walls. At the same time, because H-mode inevitably leads to the increased confinement of impurities within the plasma, which can dilute the fuel, ELMs offer an attractive mechanism for flushing out said impurities.

The filamentary nature of ELMs, that is, the tendency of a ELM to be accompanied by bursts of plasma and heat locally aligned with a magnetic field line [Kirk et al., 2006; Yun et al., 2011], and the consequences this has on the “crash” phase of an ELM event, is of particular relevance to this thesis. In recent years, the KSTAR tokamak has performed a series of experimental campaigns aimed at advancing the current understanding of ELMs. A thorough account of ELM dynamics in KSTAR can be found in Refs. [Yun et al., 2011; Thatipamula et al., 2016; Kim et al., 2018], in which three distinct stages of ELM filament evolution in

KSTAR are detailed:

1. The initial growth of the ELM filaments near the last closed flux surface (LCFS), which grow to a saturated state in $\sim 300\mu s$.
2. The interim quasi-steady state of saturated filaments typically persisting for $\sim 100\mu s$.
3. The collapse of the H-mode pedestal through multiple toroidally and poloidally localised filament bursts occurring over a time-scale of $\sim 100\mu s$.

The third and final stage of the ELM dynamics is referred to henceforth as the “ELM crash”. It is this ELM crash phase which is of relevance to this thesis, and will be revisited in Chapter 4.

1.4 Ion cyclotron emission and the Magnetoacoustic cyclotron instability

A consequence of the reaction shown in Eq. 1.2 is the production of fusion-born alpha particles. Once they are produced, the alpha particles must transfer their energy to the bulk plasma (the deuterium and tritium), in order to provide the plasma heating necessary to achieve self-sustained nuclear burning. Instabilities within the plasma that arise on time-scales smaller than the the time-scales on which the alpha particles deliver their energy to the bulk plasma through collisional processes, could either reduce the efficiency of this energy transfer, or perhaps even enhance it through “alpha-channelling” concepts [Fisch, 1995a; Herrmann and Fisch, 1997; Fisch, 1995b]. An inverse ion cyclotron emission concept for alpha-channelling was recently explored by [Cook et al., 2017]. To this end, it is imperative that diagnostics are developed to monitor the confinement and evolution of fusion born ions. One such method to do so is by exploiting the plasma instabilities themselves.

In this section I will discuss a candidate for such treatment, the experimentally observed ion cyclotron emission “ICE”, which comprises strongly suprathermal emission having spectral peaks at multiple fast ion cyclotron harmonic frequencies. ICE is a passive, non-invasive diagnostic that can be detected by the antennae used for ICRH, or a dedicated RF frequency probe [McClements et al., 2015]. I will first begin with an account of experimental ICE observations in MCF experiments and follow with a discussion of the analytical linear theory developed to interpret these observations. This section then ends with an overview of recent simulation results which turn out to be indispensable in the quest to understand ICE.

1.4.1 Ion cyclotron emission (ICE)

Before going any further, it is convenient to distinguish between the three main types of ICE observed experimentally:

1. ICE resulting from fusion products - “FP ICE”.
2. ICE resulting from neutral beam injection - “NBI ICE”.
3. ICE due to ICRF wave-accelerated minority ions (see Sec. 1.3.2) - “ICRF ICE”.

Some of the earliest accounts of ICE were seen in pure deuterium JET plasmas and are detailed in [Cottrell and Dendy, 1988; Schild et al., 1989] and references therein. In [Cottrell and Dendy, 1988], ICE was observed during both Ohmic and neutral beam heated JET plasmas. It was concluded that the peak frequency of emission corresponded to harmonics of suprathreshold fusion-born protons in the outer edge plasma. It is noted here that the large radial excursions of protons trapped on banana orbits similar to that shown in Fig. 1.5 result in a significant fraction of protons with a large perpendicular velocity in the edge region. The work of Ref. [Schild et al., 1989] detailed observations of ICE inverted sawtooth oscillations in the edge of a JET plasma. These oscillations were correlated with both a soft X-ray signal and D_α emission during the discharge, and confirmed the origin of ICE as the outer region of the plasma, paving the way for future analytical models. Around the same time, ICE was observed in TFTR [Greene and TFTR, 1990], however this ICE was shown to be the result of fast particles arising from deuterium neutral beam injection - NBI ICE, as opposed to the fusion driven FP ICE observed in JET.

It was not long before FP ICE from alpha particles in deuterium-tritium (DT) reactions was observed on both devices [Cottrell et al., 1993; Cauffman et al., 1995]. In perhaps the most significant experimental ICE result of the era, it was found that the measure neutron flux is linearly proportional to ICE intensity over six orders of magnitude [Cottrell et al., 1993]. The same work noted that the ICE signal disappeared within 0.5ms whenever there was a large amplitude ELM, due to the expulsion of energetic ions from the edge emitting region. It is worth remembering this last result, later in chapter 4, the conjunction between ELM physics and fast ion physics is explored further with the help of modern plasma diagnostics [Thatipamula et al., 2016; Kim et al., 2018]. Similar results were reported in TFTR plasmas [Cauffman et al., 1995], in which it was observed that in some TFTR discharges, the emission was brief when compared with its JET counterpart. This difference was

explained in terms of the Alfvénic nature of ICE - the TFTR alpha particles responsible for the brief ICE signal had a speed less than the local Alfvén speed, whereas the JET alpha particles had a speed greater than the local Alfvén speed. These JET DT experiments were heated both Ohmically and by NBI. Two years later, FP ICE during JET ICRH discharges was detected for the first time [McClements et al., 1999]. Figure 1.6 shows example ICE spectra from both JET and TFTR in which a sequence of peaks at multiples of the alpha cyclotron frequency can be seen.




Figure 1.6: Left: ICE intensity as a function of frequency for fusion born ICE in JET. Reproduced from [Cottrell et al., 1993]. Right: ICE intensity as a function of frequency for fusion born ICE in TFTR. Reproduced from [Cauffman et al., 1995]. In both plots, evenly spaced sequential harmonics of the alpha particle cyclotron frequency can be seen.

ICRF ICE was reported for the first time in the JET tokamak [Cottrell, 2000]. After switching on the ICRH, emission at the first harmonic of the proton cyclotron frequency was observed with a $\sim 0.4s$ time delay. The ICRH protons were accelerated to MeV energies by the applied wave giving them a distribution function, and hence orbital structure, similar to the fusion products responsible for the previously observed FP ICE. The similar nature of the ICRF ICE to the more familiar FP ICE opened up the possibility of stimulating ICE so as to study the physics experimentally.

The ICE discussed thus far was detected at the outer midplane edge of the tokamak, the first instance of FP ICE originating from the centre of the tokamak was reported on JT-60U [Kimura et al., 1998]. This ICE was also observed to correlate temporally with ELMs, in particular, the associated change in density. In Ref. [Shalashov et al., 2003], the properties of NBI ICE were investigated on the W7-AS stellarator. These W7-AS experiments used both radial (approximately perpendicular) and tangential NBI and observed ICE in both cases. It was found that radially

injected beams increased the intensity of the low ICE harmonics already present during tangential beam injection, as well as increasing the number of observed ICE harmonics. The spatial structure of ICE at the outer midplane edge due to both fusion products (tritium and helium three) and NBI (deuterium) was investigated in JT-60U by [Ichimura et al., 2008] and [Sato et al., 2010]. Curiously, FP ICE was only excited during tangential NBI, and was strongly dependent on the plasma density. More stellarator results, this time from hydrogen plasmas in LHD, were reported by [Saito et al., 2009]. The ICE had spacing equal to the proton cyclotron frequency and was synchronised with perpendicular proton NBI, suggesting that the latter might drive the ICE. Two different types of ICE were detected in LHD some years later [Saito et al., 2013]. The first, was again perpendicular NBI ICE. The second was synchronised with bursts of toroidal Alfvén eigenmodes (TAEs) [Heidbrink and Sadler, 1994; Duong et al., 1993], and thought to be the result of the bursts selectively transporting energetic ions from the core to the outer region of the plasma. Interestingly, the fast particles responsible for the NBI ICE in LHD described above do not exist on confined orbits and are lost from the plasma.

In the DIII-D tokamak, the fidelity of ICE as a fast ion diagnostic was again demonstrated when it was found to give a reliable measurement of the timing and magnitude of fast ion losses during off-axis fishbones [Heidbrink et al., 2011]. ICRF ICE was again detected in the JET using the SHAD detection system and first reported in Ref. [Jacquet et al., 2011]. The interpretation of these results is a subject of this thesis and is examined more in chapter 5. A thorough account of recent ICE measurements on ASDEX Upgrade (AUG) was carried out by [D’Inca, 2014] in which all three types of ICE outlined above were observed, with FP ICE being observed both the outer-midplane edge and the core plasma. Recent measurements of NBI ICE on DIII-D were reported in [Pace et al., 2016] and were found to correlate with the losses of neutral beam injected ions, which again further strengthens the case for an ICE diagnostic in future fusion devices. Observations of ICE in the spherical tokamak NSTX-U were reported in [Fredrickson et al., 2017]. This ICE signal was assumed to originate from deuteron NBI, and its frequency is such that appears to originate from within the bulk plasma, at approximately half minor radius.

Recently, both NBI and FP ICE have been observed in the KSTAR tokamak [Thatipamula et al., 2016; Kim et al., 2018]. It is these KSTAR ICE observations that are the dominant focus of this thesis, and they are discussed in great detail in chapter 4. Core FP ICE with spacings equal to both the local proton, deuteron, and tritium cyclotron frequencies has recently been observed in AUG deuterium plasmas [Ochoukov et al., 2018, 2019]. This ICE was commonly observed during

the NBI start up phase, and also under steady state conditions, persisting for more than one second. Recent DIII-D experiments have observed core ICE from the 2nd, 3rd, and 4th harmonics of the deuteron cyclotron frequency, which is believed to be NBI ICE, with FP ICE possibly contributing to the overall ICE signal [Thome et al., 2018]. Finally, ICE has recently been observed in both hydrogen and deuterium NBI heated plasmas in the TUMAN-3M device [Askinazi et al., 2018]. The fusion reactivity is low in the hydrogen plasmas, and if all the fusion products are unconfined in the deuterium plasmas, this would exclude the possibility of FP ICE. Multiple ICE signals with frequency spacing equal to the proton and deuteron cyclotron frequencies were observed in both the high and low field sides of the torus, as well as in the core. The physical origin of these ICE signals is unclear, but it seems reasonable to suppose it is NBI ICE.

Table 1.3 summarises the experimentally observed ICE discussed above. For each entry the nature of the ICE - FP, NBI or ICRF - is given, along with a reference and year of publication (not necessarily the year in which the ICE was first detected). Asterisks denote ICE that was observed at a location other than, or in addition to, the plasma edge. This list is by no means comprehensive, but is intended to provide an approximate timeline of the experimental research related to ICE, hopefully providing more context to the rest of this thesis.

Table 1.3: Timeline of experimental ICE observations in MCF plasmas.

Device	Nature of ICE	Year	Reference
JET	FP	1988	[Cottrell and Dendy, 1988]
JET	FP	1989	[Schild et al., 1989]
TFTR	NBI	1990	[Greene and TFTR, 1990]
JET	FP	1993	[Cottrell et al., 1993]
TFTR	FP	1995	[Cauffman et al., 1995]
JT-60U	FP*	1998	[Kimura et al., 1998]
JET	FP	1999	[McClements et al., 1999]
JET	ICRF	2000	[Cottrell, 2000]
W7-AS	NBI	2003	[Shalashov et al., 2003]
JT-60U	NBI	2008	[Ichimura et al., 2008]
LHD	NBI	2009	[Saito et al., 2009]
JT-60U	NBI+FP	2010	[Sato et al., 2010]
DIII-D	? ¹	2011	[Heidbrink et al., 2011]
JET	ICRF	2011	[Jacquet et al., 2011]
LHD	NBI+TAE ²	2013	[Saito et al., 2013]
ASDEX U	NBI+FP*+ICRF	2014	[D’Inca, 2014]
DIII-D	NBI	2016	[Pace et al., 2016]
KSTAR	NBI+FP	2016	[Thatipamula et al., 2016]
NSTX-U	NBI*	2017	[Fredrickson et al., 2017]
DIII-D	NBI*+FP* ³	2018	[Thome et al., 2018]
TUMAN-3M	NBI* ⁴	2018	[Askinazi et al., 2018]
ASDEX U	NBI*	2018	[Ochoukov et al., 2019]

¹The origin of the emission is unclear. Refer to the original manuscript for further details.

²This is not one of the three main types of ICE defined at the beginning of this section. It is thought to be due to the motion of particles from the core of the plasma to the edge resulting from TAE bursts.

³FP ICE may contribute to the overall ICE signal.

⁴The origin of the emission is not completely clear.

It is worth noting that ICE is an ubiquitous plasma phenomenon, and not just observed in magnetic fusion devices. ICE has been detected in solar-terrestrial plasmas [McClements and Dendy, 1993; Dendy and McClements, 1993; McClements et al., 1994; Dendy et al., 2002], and is possibly present downstream of supernova remnant shocks [Rekaa et al., 2014]. A contemporary example of ICE observed in the Earth’s Van Allen belts is shown in [Posch et al., 2015]. Figure 4a from this paper has been reproduced below. Plotted is a spectrogram of the y-component of the electric field in the radiation belt as the Van Allen probes spacecraft executed an eccentric orbit through a region of varying magnetic field. The vertical axis plots the frequency of excited waves while the horizontal axis plots the time during this particular orbit. Excitation of compressional waves at several low hydrogen cyclotron harmonics of hydrogen is apparent.




Figure 1.7: Spectrogram of E_y component of the field measurement by the Van Allen Probes spacecraft. Reproduced from [Posch et al., 2015].

1.4.2 The Magnetoacoustic Ion cyclotron instability (MCI)

The emission mechanism behind ICE is thought to be the magnetoacoustic cyclotron instability (MCI) which was originally formulated theoretically by [Belikov and Kolesnichenko, 1976]. The MCI occurs when a minority fast ion population enters into cyclotron resonance with a fast Alfvén wave supported by the bulk plasma propagating nearly perpendicular to the background magnetic field. As the frequencies of the MCI lie along the fast Alfvén branch, for purely perpendicular propagation it is therefore a transverse wave which is polarised in the binormal (poloidal) direction. The distribution of collectively radiated energy between different cyclotron harmonics depends on the character of the MCI. A key dimensionless parameter which determines the threshold for instability is the ratio of energetic ion perpendicular velocity to the Alfvén speed, which can be inferred from both linear theory and fully nonlinear self-consistent PIC simulations. Motivated by the plethora of

ICE measurements in the early 90's, the development of the theory was catapulted forward. The work of [Dendy et al., 1992] under the local, straight magnetic field, homogeneous plasma approximation captured many of the features of the early JET and TFTR experiments and is applicable for scenarios in which the drift motion of the ions across magnetic field lines is slow compared to the growth rate of the instability - “rapid instability”. For this first work, an isotropic, infinitesimally thin shell distribution was adopted for the energetic fusion born ions. This was subsequently extended in Ref. [Dendy et al., 1993] to shell distributions of finite thickness. The dispersion relation of the MCI under these conditions can be found in Eq. 31 of Ref. [Dendy et al., 1993]. The theory of the MCI was later extended further to include waves with finite parallel wavenumber, under the approximation that the electric field is polarised in the plane perpendicular to the magnetic field, thereby excluding Landau damping effects, this time using a drifting ring beam distribution for the energetic ions [Dendy et al., 1994a]. The inclusion of this effect is especially important, because the resulting Doppler shift allows for wave growth even in the presence of majority ion cyclotron damping. The preceding works were motivated by observations of super-Alfvénic FP ICE in JET. This expression was analysed in the context of sub-Alfvénic FP ICE observed in TFTR [McClements et al., 1996]. In addition, another emission mechanism was proposed to explain the sub-Alfvénic NBI ICE observed in TFTR, in which electrostatic ion cyclotron harmonic waves can be destabilised by the minority NBI ions [Dendy et al., 1994b, 1995].

The effect of toroidal drift effects on the MCI was investigated in Refs. [Fülöp et al., 1997; Fülöp and Lisak, 1998], in which the authors found a significant enhancement to the instability growth rate. In the “weak instability” limit, in which the drift motion of ions is fast compared to the growth rate, theories involving the excitation of Compressional Alfvén eigenmodes (CAEs) have been developed which include the associated toroidal eigenmode structure [Fülöp et al., 2000; Gorelenkov and Cheng, 1995; Gorelenkov, 2016; Gorelenkov et al., 2014; Kolesnichenko et al., 1998, 2000; Smith et al., 2003; Hellsten and Laxaback, 2003; Hellsten et al., 2006; Smith and Verwichte, 2009]. Recently, linear analysis of ICE due to the MCI has successfully been applied to JT-60U [Sumida et al., 2017], along with a detailed study of the characteristics of the energetic ion velocity distribution responsible for driving the ICE [Sumida et al., 2018].

In this thesis the instability is considered under the local, straight field, homogeneous plasma approximation. By considering the instability in this locally uniform or slab-like limit, it becomes computationally practical to carry out fully

nonlinear, fully kinetic, and fully self-consistent simulations of the MCI. In addition, this approach enables benchmarking of computational results at early times in the simulations against analytical linear theory. The nonlinear effects that are accessible only through direct numerical simulation turn out to be very important for the understanding of ICE. After all, the ICE spectra observed in fusion devices, and indeed elsewhere, are likely to represent the saturated nonlinear regime.

Recent advances in computational physics have enabled further study of the MCI and ICE. PIC simulations, described in [Cook et al., 2013], carried the MCI instability into its non-linear regime for the first time. Using the PIC code EPOCH, fully kinetic ions and electrons, together with self-consistent electric and magnetic fields were modelled for plasma parameters aligned relevant to JET conditions. These PIC results were reinforced by subsequent ‘‘Hybrid’’ simulations of the MCI over longer physical time scales, facilitating further exploration of the nonlinear regime of the MCI [Carbajal et al., 2014; Carbajal, 2015]. In the hybrid approximation, ions are treated as particles and electrons as a massless neutralising fluid [Winske et al., 2003]. Preliminary work has begun on a hybrid version of EPOCH, which borrows many ideas from the code used in the work of [Carbajal, 2015]. Details of this can be found in Appendix A. More recently, [Cook et al., 2017] investigated a stimulated emission counterpart to the MCI under JET-like conditions. An applied wave resulted in the extraction of energy from a minority alpha particle population to a majority deuteron population, ‘‘alpha particle channelling’’. The pioneering result of [Cottrell et al., 1993], which showed that the intensity of ICE in JET was found to scale linearly with neutron flux and minority ion concentration, found its simulation counterpart in the work of [Carbajal et al., 2017], in which the observed scaling was reproduced using hybrid-PIC simulations. Finally, the first nonlinear MCI simulations of NBI proton driven ICE in the LHD stellarator have been successful in reproducing the observed spectrum [Reman et al., 2016; Dendy et al., 2017].

In the analytical formalism of the MCI, the minority energetic ion particle population is usually modelled as a drifting ring beam distribution in velocity space according to

$$f(v_{\parallel}, v_{\perp}) \propto \exp\left(-\frac{(v_{\parallel} - v_d)^2}{v_r^2}\right) \exp\left(-\frac{(v_{\perp} - u_0)^2}{u_r^2}\right),, \quad (1.22)$$

where v_{\perp} and v_{\parallel} are the minority ion velocity components perpendicular and parallel to the background magnetic field respectively. The magnitude of the initial perpendicular velocity and initial drift along the background magnetic field are de-

noted by u_0 and v_d respectively, and the perpendicular and parallel velocity spreads are given by u_r and v_r respectively. Variants of this distribution function are used throughout the simulations described in the preceding paragraph and those of this thesis. Without a velocity spread in the perpendicular direction Eq. 1.22 reduces to

$$f(v_{\parallel}, v_{\perp}) \propto \exp\left(-\frac{(v_{\parallel} - v_d)^2}{v_r^2}\right) \delta(v_{\perp} - u_0). \quad (1.23)$$

This simplifies further in the case of no parallel spread

$$f(v_{\parallel}, v_{\perp}) \propto \delta(v_{\parallel} - v_d) \delta(v_{\perp} - u_0). \quad (1.24)$$

Finally, in the case of no parallel spread, no perpendicular spread, and no parallel drift, the distribution is given by

$$f(v_{\perp}) \propto \delta(v_{\parallel}) \delta(v_{\perp} - u_0). \quad (1.25)$$

It is helpful to view the structure of these distribution functions in 3D phase space, as shown in Fig. 1.8.

The distribution function given in Eq. 1.25 is one predominantly used in previous simulations of the MCI as well as chapters 3, 4, 5, and 7 of this thesis. In chapter 6 we show results from simulations in which the minority energetic ions are initialised as a ring beam distribution with zero parallel velocity and a finite perpendicular velocity spread, according to

$$f(v_{\parallel}, v_{\perp}) \propto \exp\left(\frac{-(v_{\perp} - u_0)^2}{u_r^2}\right) \delta(v_{\parallel}); \quad (1.26)$$

along with simulations in which the minority energetic ions are initialised according to a spherical shell distribution with

$$f(v) \propto \exp\left(\frac{-(v - v_0)^2}{v_T^2}\right). \quad (1.27)$$

Here v is the magnitude of the total velocity vector, v_0 is the magnitude of the initial velocity, and v_T is the velocity spread, corresponding to the thickness of the shell in velocity space.




Figure 1.8: Phase space plots showing the variants of the distribution function used to represent the minority ion species in MCI simulations. All distributions have identical values of the initial perpendicular velocity u_0 . The parallel drift velocity v_d , parallel spread v_r , and perpendicular spread u_r are expressed as multiples of u_0 . Top left: Corresponds to Eq. 1.22. Top right: Corresponds to Eq. 1.23. Bottom left: Corresponds to Eq. 1.24. Bottom right: Corresponds to Eq. 1.25.

Chapter 2

Computational methods: Particle-In-Cell codes and higher order spectra

The following chapter will give an overview of the particle-in-cell (PIC) simulation technique used in this thesis, as well as bicoherence analysis, a higher order spectral technique used to obtain some of the results in chapters 4, 6, and 5.

2.1 Particle-In-Cell codes

Particle in cell (PIC) codes, provide a basis with which to fully resolve the kinetic profiles of all particles. The plasma is represented by a distribution function f of plasma species, and the evolution of the positions and velocities of these particles is governed according to the Maxwell-Lorentz system of equations described in section 1.2.

In order to make simulations computationally tractable, pseudo/macro particles are used to represent collections of real plasma particles, each computational particle contributes a small amount Δf to the overall probability density function of the plasma species. The physical domain is defined by a grid, with the electromagnetic fields being specified at each grid cell, as opposed to at the location of every particle. Such macro particles have the same charge to mass ratio (and hence the same orbit) as real particles, but much larger values of charge and mass. As they represent an ensemble of real particles, they occupy a finite volume and have a shape, which is unchanged under normal motion. These macro particles are free to move from the grid, often traversing several grid cells.

To put this into context, a typical tokamak simulation may have a physical particle density of 10^{20}m^{-3} , but, in a PIC simulation, this will be represented by an ensemble of macro particles, having a number density of say, 10^{12}m^{-3} . Thus, each macro particle represents 10^8 physical particles, carrying 10^8 times as much current and charge density as their physical counterparts. At this point one may be concerned with the problem of self-consistently evolving the electromagnetic field equations, this is achieved by using the charge and current of the macro-particles, which are interpolated from particles positions to their nearest grid points, as source terms in Gauss’s and Ampere’s law.

The “EPOCH” PIC code is used for the simulations described in the following chapters. It follows an explicit integration scheme, is fully relativistic, second order accurate and MPI parallelised [Arber et al., 2015]. We first start with a statement of the PIC procedure. To advance the system from a time level n to a time level $n+1$ we:

1. Update the electromagnetic fields by half a timestep

$$\begin{aligned}\mathbf{E}^n &\rightarrow \mathbf{E}^{n+1/2} \\ \mathbf{B}^n &\rightarrow \mathbf{B}^{n+1/2}.\end{aligned}$$

2. Update the particle positions and velocities by a full timestep

$$\begin{aligned}\mathbf{v}^n &\rightarrow \mathbf{v}^{n+1} \\ \mathbf{x}^n &\rightarrow \mathbf{x}^{n+1}.\end{aligned}$$

3. Update the particle current by a full timestep

$$\mathbf{J}^n \rightarrow \mathbf{J}^{n+1}.$$

4. Update the electromagnetic fields by another half a timestep

$$\begin{aligned}\mathbf{E}^{n+1/2} &\rightarrow \mathbf{E}^{n+1} \\ \mathbf{B}^{n+1/2} &\rightarrow \mathbf{B}^{n+1}.\end{aligned}$$

In practice steps 2 and 3 are performed concurrently in the same block of code, but have been separated here for clarity. Each of the above steps will now be explained in more detail.

2.1.1 Field update

The electromagnetic fields in a PIC code are updated using Faraday’s law for the magnetic field and Ampere’s law for the electric field. The electromagnetic fields are defined on a Yee staggered grid [Yee, 1966], which simplifies the calculation of the field derivatives to second order accuracy. The use of the Yee grid also means provided $\nabla \cdot \mathbf{B}$ is initially 0, it remains so throughout the simulation. The current \mathbf{J} is calculated using the Villasenour-Buneman method [Villasenor and Buneman, 1992], which means Gauss’s law is always satisfied, this is discussed in more detail in section 2.1.4. This means only two of the four Maxwell’s equations are needed to evolve the system. Figure 2.1 shows how the fields are defined on the Yee grid in one spatial dimension. The quantities B_y, B_z , and E_x are defined on the cell edge, so that their spatial derivative in the x-direction is automatically defined at the correct location. For instance, $B_y \propto \partial_x E_z$, and the derivative of E_z is automatically defined at the location of B_y . Note, the derivatives in the y and z directions are always zero in a 1D PIC code. In a 1D scheme, the components of current are defined at the same locations as the corresponding electric field component to enable easy computation of the discrete form of Ampere’s law. Quantities such as density and temperature, “bulk” quantities, are defined at the cell centre. Figure 2.1 shows schematically the macro particles on the grid, occupying a finite volume of the grid cell, with velocity vectors, denoted by arrows of varying length, pointing in arbitrary directions - whilst the code has only one spatial dimension, the full phase space is represented. The discretised field equations are

$$\mathbf{E}^{n+1/2} - \mathbf{E}^n = \frac{\Delta t}{2} \left(c^2 \nabla \times \mathbf{B}^n - \frac{\mathbf{J}^n}{\epsilon_0} \right), \quad (2.1)$$

$$\mathbf{B}^{n+1/2} - \mathbf{B}^n = -\frac{\Delta t}{2} \left(\nabla \times \mathbf{E}^{n+1/2} \right). \quad (2.2)$$

At this point the particle position and velocity are updated to find \mathbf{J}^{n+1} .

$$\mathbf{B}^{n+1} - \mathbf{B}^{n+1/2} = -\frac{\Delta t}{2} \left(\nabla \times \mathbf{E}^{n+1/2} \right), \quad (2.3)$$

$$\mathbf{E}^{n+1} - \mathbf{E}^{n+1/2} = \frac{\Delta t}{2} \left(c^2 \nabla \times \mathbf{B}^{n+1} - \frac{\mathbf{J}^{n+1}}{\epsilon_0} \right). \quad (2.4)$$




Figure 2.1: Schematic of the 1D Yee grid used in the 1D3V version of the EPOCH PIC code.

2.1.2 Particle Push

The relativistic particle trajectories are modelled along the grid using an explicit finite difference method known as the Boris scheme. Other finite difference schemes such as the Leapfrog method can be used, but because of its long term accuracy, the Boris method has become the de-facto standard for PIC codes. Such codes rely on interpolating the electromagnetic fields from the grid to the particle position to provide the Lorentz force. Once this force is known, the particle velocities and positions are updated from time n to time $n+1$, this is known as the particle push [Birdsall and Langdon, 1985]. It should be noted that for highly relativistic plasmas, the Boris push can lead to a spurious force, and an alternative “Vay” push has been proposed to remedy this [Vay, 2008]. The plasmas simulated in this thesis are non-relativistic, meaning the Boris push still works extremely well. For a given particle, the Boris method is as follows.

First, the particle position is updated by half a time-step from time $t = n\Delta t$ to $t = (n + 1/2)\Delta t$

$$\mathbf{x}^{n+1/2} = \mathbf{x}^n + \frac{\Delta t}{2} \mathbf{v}^n. \quad (2.5)$$

The relativistic Lorentz force law is discretised according to

$$\frac{\mathbf{u}^{n+1} - \mathbf{u}^n}{\Delta t} = \frac{q}{m} \left(\mathbf{E}^{n+1/2} + \frac{\mathbf{u}^{n+1/2}}{\gamma^{n+1/2}} \times \mathbf{B}^{n+1/2} \right). \quad (2.6)$$

where the relativistically correct velocity is $\mathbf{u} = \gamma \mathbf{v}$. The relativistic factor is given by $\gamma = \left[\left(\frac{|\mathbf{u}|}{c} \right)^2 + 1 \right]^{1/2}$ where m is the rest mass. As γ is a function of \mathbf{u} , it changes with time and one must be careful to update it properly. Equation 2.6 presents a problem as the velocity at the half time step is not known. To overcome this, it is

set to be an average of the velocities at the two known time levels

$$\frac{\mathbf{u}^{n+1} - \mathbf{u}^n}{\Delta t} = \frac{q}{m} \left(\mathbf{E}^{n+1/2} + \frac{(\mathbf{u}^{n+1} + \mathbf{u}^n)}{\gamma^{n+1/2}} \times \mathbf{B}^{n+1/2} \right). \quad (2.7)$$

This, however, presents another problem, the velocity we are trying to find, appears on both the left and right side of the equation and is “tangled up” with $E^{n+1/2}$. To handle this, we use the Boris rotation algorithm. This splits Eq. 2.6 into two parts [Birdsall and Langdon, 1985]. The first is responsible for the acceleration due to the electric field. We define

$$\mathbf{u}^n = \mathbf{u}^- - \frac{q\Delta t}{2m} \mathbf{E}^{n+1/2}, \quad (2.8)$$

$$\mathbf{u}^{n+1} = \mathbf{u}^+ + \frac{q\Delta t}{2m} \mathbf{E}^{n+1/2}. \quad (2.9)$$

These can substitute the following equations into Eq. 2.9 to give

$$\frac{\mathbf{u}^+ - \mathbf{u}^-}{\Delta t} = \frac{q}{2m\gamma^{n+1/2}} (\mathbf{u}^+ + \mathbf{u}^-) \times \mathbf{B}^{n+1/2}. \quad (2.10)$$

Note that the term $\mathbf{u}^+ + \mathbf{u}^-$ on the right side of Eq. 2.10 can be replaced with $\mathbf{u}^{n+1} + \mathbf{u}^n$. The above is thus a rotation of $\mathbf{u}^+ + \mathbf{u}^-$ about an axis parallel to \mathbf{B} . Notice how in the above equations γ has remained defined at time level $n+1/2$. It can be shown that γ is constant throughout the step and we can say $\gamma^{n+1/2} = \gamma^- = \gamma(\mathbf{u}^-)$.

The procedure for the rotation is now:

1. Add half the electric field acceleration to \mathbf{u}^n using Eq. 2.8 to obtain \mathbf{u}^- .
2. Apply the rotation according to 2.10 to get \mathbf{u}^+ .
3. Add the remaining half of the electric field acceleration to \mathbf{u}^+ using Eq. 2.9 to obtain \mathbf{u}^{n+1} .

Finally, now that the velocity has been updated, we can advance the particles position by a further half timestep according to

$$\mathbf{x}^{n+1} = \mathbf{x}^{n+1/2} + \frac{\Delta t}{2} \mathbf{v}^{n+1} = \mathbf{x}^{n+1/2} + \frac{\Delta t}{2\gamma^{n+1/2}} \mathbf{u}^{n+1} \quad (2.11)$$

This yields a particle push which is second order accurate in time. The current \mathbf{J} is updated using the the Villasenor and Buneman current calculating scheme

[Villasenor and Buneman, 1992], by solving the following additional equation

$$\frac{\partial \rho}{\partial t} = \nabla \cdot \mathbf{J}. \quad (2.12)$$

At the end of the particle push, once the velocity and position are defined at time level $n + 1$, the particle position is updated by a further half timestep, to time level $n + 3/2$. This is used with the particle position at time level $n + 1/2$, which was stored previously at the beginning of the particle push, to obtain a time centred evaluation of $\partial \rho / \partial t$, which is then used to update \mathbf{J}^n to \mathbf{J}^{n+1} . The current update is thus second order accurate in time, but one sided space [Esirkepov, 2001]. For spatial directions not captured by the PIC code, e.g. y and z in the 1D3V version of EPOCH, the current is obtained by taking moments of the particles, this is detailed in the next section.

2.1.3 Particle shape functions

The finite volume of macro-particles, which may be larger than the volume of a grid cell, has implications when calculating the electric and magnetic fields at the location of the particles during the particle push. Typically, particle weighting functions are used to interpolate the fields from the grid to the particles. The higher the order of the weight function, the smoother the fields are at the particles location, and the less noisy the final solution is. A simple, first order, weighting function is one in which the particle fills the cell uniformly, a top-hat function. A better approach would be to assume the particle occupies a triangular shape in configuration space. The peak of the triangle is located at the position of the centre of mass of the macro particle, and the triangle has a width $2\Delta x$, twice the grid cell width Δx . This second order triangular weighting function requires a summation over three cells, the cell in which the particle lies, and its two nearest neighbours. Higher order shape functions using splines to interpolate the fields to the particles position also exist. EPOCH is able to use third order spline interpolation, which involves a sum over five cells. In general, higher order particle shape functions yield better results, but with an increased computational cost.

A particle that lies somewhere in cell j experiences contributions to its own field from the fields surrounding it which are defined at fixed points on the grid. For a given weighting function $S_m(x)$ of order m , the field F_i of a particle at position x_i , where x_i lies in the cell centred at x_j , and this cell has a field F_j which is constant

over the cell, is given by

$$F_i = \frac{1}{\Delta x} \sum_j F_j \int_{x_j - \frac{\Delta x}{2}}^{x_j + \frac{\Delta x}{2}} S_m(x_i - x) dx. \quad (2.13)$$

Where the range of the summation depends on the order of the weighting function. The triangular shape function, shown in Fig. 2.2 along with the tophat shape function, is given by

$$S_1(x_i - x) = \begin{cases} 1 - \frac{|x_i - x|}{\Delta x} & , |x_i - x| \leq \Delta x \\ 0 & , otherwise. \end{cases} \quad (2.14)$$

This gives the total field experienced by the particle as

$$F_i = \frac{1}{2} F_{j-1} \left(\frac{1}{2} + \frac{x_j - x_i}{\Delta x} \right)^2 + F_j \left(\frac{3}{4} - \frac{(x_j - x_i)^2}{(\Delta x)^2} \right) + F_{j+1} \left(\frac{1}{2} - \frac{x_j - x_i}{\Delta x} \right)^2 \quad (2.15)$$

For the third order splines used in EPOCH, these three terms have a different form and there are an additional two terms corresponding to F_{j-2} and F_{j+2} . The same procedure is used to extract particle quantities such as density and current to the grid. With the exception of some of the resolution testing simulations discussed in chapter 3, all the PIC results shown in this thesis use third order splines.




Figure 2.2: Schematic of the first order tophat (blue) and second order triangular (red) shape functions used in the EPOCH PIC code. This figure has been reproduced from page 21 of the most recent EPOCH developers manual (at the time of writing).

2.1.4 Practical and computational considerations

Due to the nature of explicit finite difference schemes, the time-step must be chosen such that the CFL condition is satisfied in order to maintain stability. The CFL

condition in a 3D PIC code is:

$$\Delta t < \frac{((\Delta x)^2 + (\Delta y)^2 + (\Delta z)^2)^{1/2}}{c}, \quad (2.16)$$

where Δx etc are the grid spacings in their respective grid directions, and c is the speed of light. In this work, only one dimension is considered, so Eq. 2.16 can be simplified. In practice, EPOCH implements this condition automatically, but it is the responsibility of the user to choose an appropriate grid spacing. A further time step restriction arises from the need to resolve the plasma frequency ω_{pe} . EPOCH selects the most restrictive of these timesteps, which in an MCF plasma, is usually that specified by the CFL condition. The user just needs to make sure it is small enough to resolve the physics of interest. Another serious issue one must consider when running a PIC code is the resolution of the Debye length λ_D within a warm plasma

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_b T_e}{n_e e^2}}. \quad (2.17)$$

If λ_D is not resolved by the grid spacing, the code will exponentially heat, increasing the Debye length until it equals the width of the cell Δx (which is specified by the user). This exponential heating is due to the acceleration of particles by a restoring electric field which is produced by charge density fluctuations resulting from the thermal noise. The relative change to a particles velocity is larger for plasma particle with lower thermal speeds, and can cause serious problems, particularly if the growth rate of the physics under study is large, because the exponential heating may dominate the plasma dynamics in the initial stage of the simulation. For simulations using periodic boundary conditions, the resolution between successive modes in k -space is $\Delta k = \pi/L$. If the modes of interest in a simulation are very close together in wavenumber space, one requires a small Δk , hence a large simulation domain L , so as to resolve the relevant physics. For a given L , to resolve λ_D we require the number of grid cells n_x to be at least L/λ_D . For a plasma such as that found in the centre of a tokamak, where n_e is large, n_x can therefore be very large. This restriction is even more serious when we also consider the CFL condition. The timestep is directly proportional to Δx in a 1D PIC code, meaning it must decrease as Δx decreases. So for a given grid length L , if we can justify increasing the cell size by a factor two, we may get a factor four speed up in the computation of our problem. As we need to simulate less cells, and our timestep automatically doubles (assuming the CFL timestep is the most restrictive). It is easy to see why it may be tempting to push the cell size to be as large as possible, or conversely, reduce the plasma density

as much as possible. In practice, this effect can be ameliorated by using higher order particle shape functions (at increased computational cost), and/or smoothing the current \mathbf{J} in the simulation [Arber et al., 2015]. Another consideration is the number of particles per cell. The finite size and the number of macro particles means there is noise in the electro-magnetic fields. This noise acts to heat the plasma, and must therefore be kept to a minimum. The noise scales as $N^{-1/2}$, where N is the number of macro-particles per cell. Ideally, the number of macro-particles per cell would be equal to the number of physical particles, however this is not tractable. Doubling the number of particles per cell significantly increases the computational cost, but only reduces the noise in the simulation by a factor of $\sqrt{2}$. It is possible to reduce the noise levels by using higher order particle shape-functions, and/or by using the delta-f method, in which only changes in an equilibrium particle distribution are considered, as opposed to the entire distribution function. [Sydora, 1999].

There are certain practical limitations one must consider when running a PIC simulation. It is unfeasible to write all available simulation data to disk at each time-step. To this end, data is written to disk periodically in “dumps”, it is the responsibility of the user to choose this period such that the physics can be resolved. If a simulation runs for say, ten ion gyroperiods, and we want to resolve 100 ion cyclotron frequencies in a temporal Fourier transform, then, because of the Nyquist theorem, there must be at least 2000 data dumps equally spaced in time. In general, the time between output dumps Δt_{out} will never be an exact multiple of the simulation timestep Δt_{sim} . Output dump 1 may correspond to time $t = 10\Delta t_{sim}$ in the simulation, whereas dump 2 is at time $t = 21\Delta t_{sim}$, not $t = 20\Delta t_{sim}$ like we had hoped. If we want to take temporal Fourier transforms, this can be problematic, leading to spurious artefacts because not only are we approximating our signal as a periodic, but our “dt” in our Fourier transform integral is effectively changing. This problem becomes less of an issue if $\Delta t_{sim} \ll \Delta t_{out}$, which can result in an odd situation in which the user needs to force the timestep to be artificially small, increasing their computational expense significantly, just so they can get a reliable Fourier transform. This rarely happens in practice, and if it does, it is better to take a weighted average of data from multiple simulation timesteps and output this to disk. EPOCH does not currently have the functionality to do a weighted average of outputs, but it can do regular averages, which is usually sufficient. Aliasing is another important issue that must be considered before running a simulation. It is helpful to think of the 2D discrete Fourier transform as a box of frequencies and wavenumbers with finite area. The height of the box is the maximum resolvable frequency, and the width of the box is the maximum resolvable wavenumber. If

there is a wave in the simulation propagating with a large, positive group velocity, and this wave has no cut-off frequency, it will hit the top of the box and reflect back down into a region of lower frequency and higher wavenumber. Even if this wave is of no interest to us, it may then interact with some of the higher k modes, producing non-physical spectrally dense regions modes in the dispersion relation. It is then possible for said modes to further interact with other modes in the simulation. Before we know it, there is non-physical Fourier power everywhere. Data averaging between output dumps can help with this, or we can just make our box taller, by resolving a higher frequency, in the hope that when our uninteresting wave bounces back down, it does so in a region of k space of no interest to us. Even if the data is only written periodically and not at every time-step, the amount of disk space used when outputting particle data is often too much. The usual procedure is to write out the values of the field variables every chosen user time period, and to write out the particle quantities less frequently.

With all the above in mind, it is clear that to run PIC codes requires a kind of balancing act, between being able to accurately simulate the problem of interest, and being able to do so in a feasible amount of time, such is the case with most computational physics problems. The challenges associated with these problems are discussed in chapter 3.

2.1.5 Cold plasma dispersion relation

The thesis' of [Cook, 2011],[Gingell, 2013] and [Irvine, 2018] show multiple EPOCH benchmarking simulations for a wide range of plasma phenomena including electron Bernstein waves, the two stream instability, and waves in a thermal plasma. As such, here only benchmarks of a thermal plasma are shown and discussed. Figure 2.3 shows the results of two EPOCH simulations of a thermal deuterium plasma in which the background magnetic field was aligned along the z -axis perpendicular to the simulation domain (left) and along the x -axis parallel to the simulation domain (right). The parameters were chosen to be representative of the JET tokamak, with a magnetic field $|\mathbf{B}| = B = 2T$, density $n_e = 1 \times 10^{19} \text{m}^{-3}$, and $T_e = T_D = 1\text{keV}$. The simulation domain was chosen to be 2m, and the cell size chosen to be 0.95 times the Debye length, meaning 28,500 cells were needed. The number of particles per cell was set to 100 and both simulations used periodic boundary conditions and ran for 20 deuteron gyro-periods. In each panel, the cold plasma dispersion relation is over plotted in black and is seen to closely match the warm plasma result produced by the PIC simulation. In both cases, the resolution in frequency and wavenumber space is the same, but the right panel, corresponding to parallel

wave propagation, is not well resolved due to the sharply increasing gradient of the whistler wave. This is perhaps not an issue for this demonstration of EPOCH, but if say, we wanted to impose an extra wave in our simulation, with a frequency around $14\omega_{cp}$, and investigate how it interacts with the main whistler branch; we would miss this interaction, because our Δk is so large that we have effectively “skipped” this mode in our simulation, as is evidenced by the large gap between $\sim 13\omega_{cp}$ and $\sim 15\omega_{cp}$. It is therefore recommended to calculate Δk in advance when running a PIC simulation with periodic boundary conditions to make sure all the modes of interest are captured.




Figure 2.3: Warm plasma dispersion relation as simulated by EPOCH. Left: Perpendicularly propagating waves. Right: Parallel propagating waves. Shading indicates the spectral power in B_z field component and is plotted on a log10 scale.

2.2 Bicoherence and Bispectral analysis

Significant portions of this thesis examine nonlinear wave coupling in both simulations and experiment. One method with which to quantify this coupling is by using bispectral analysis. There is extensive literature which describes higher order spectral techniques including bispectral analysis. For general information see [de Witt, 2003; Kim et al., 1980; Kaup et al., 1979; Kravtchenko-Berejnoi et al., 1995] and for plasma-specific applications see, for example, Holland et al. [2002]; Moyer et al. [2001]; Yamada et al. [2008]; Itoh et al. [2017]. For a KSTAR plasma specific application, directly relevant to the work in this thesis, see [Lee et al., 2016]. An early account of the application of higher order spectral techniques to plasma physics is given in Ref. [Kim and Powers, 1979]. Bispectral analysis has been previously been successfully applied to the MCI [Carbajal et al., 2014].

Any three waves interacting nonlinearly must satisfy, to good approximation, the frequency and wavenumber matching conditions

$$f_3 = f_1 + f_2,$$

$$\mathbf{k}_3 = \mathbf{k}_1 + \mathbf{k}_2,$$

where waves “1” and “2” interact to produce wave “3”. To measure the amount of phase coherence between three modes that obey the above resonance conditions, one can compute the bispectrum. Defining $F(f_1)$ as the complex Fourier transform of a quantity (for instance an electromagnetic field component) at frequency $f = f_1$, and $F^*(f_1)$ as its conjugate, the bispectrum is defined as

$$b_s^2(f_1, f_2) = |\langle F(f_1) F(f_2) F^*(f_1 + f_2) \rangle|^2, \quad (2.18)$$

where the brackets $\langle \cdot \rangle$ denote averaging over time. One can normalise the bispectrum to obtain the bicoherence. This can be done in several ways [de Witt, 2003; Kravtchenko-Berejnoi et al., 1995], one of which is to use Schwartz’s inequality

$$b_c^2(f_1, f_2) = \frac{|\langle F(f_1) F(f_2) F^*(f_1 + f_2) \rangle|^2}{\langle |F(f_1) F(f_2)|^2 \rangle \langle |F^*(f_1 + f_2)|^2 \rangle}. \quad (2.19)$$

In practice, an ensemble average can replace the average over time. One thus computes several successive Fourier transforms of the same signal, sliding the Fourier transform window along the signal as we do so. These windows can overlap to some extent, which allows us to obtain improved spectral resolution, provided one is careful not to induce correlation where there is none. The number of *independent* Fourier transforms M , must be large enough so that the value of the bicoherence b_c is statistically significant. For significant coupling the variance of b_c is given by [Kim and Powers, 1979]

$$\text{Var}(b_c) \simeq \frac{1}{M} (1 - b_c^2). \quad (2.20)$$

Some accounts differ slightly in that there is a factor 4 in the numerator of the right hand side of Eq. 2.20. Regardless of this factor, we can safely say that if $b_c > 1/M$, we have a statistically significant result. All bicoherence dependant conclusions in this thesis satisfy this condition.

The bispectrum/bicoherence can either be an “auto” quantity, in which interactions between waves in one signal are computed, or a “cross” quantity, in which interactions between waves from two different signals are considered. The bispectrum, Eq. 2.18, measures the extent of phase coherence due to the nonlinear coupling

between three waves that satisfy the frequency and wavenumber matching criteria above. The bicoherence, Eq. 2.19, is a normalised bispectrum bounded between 0 and 1, and quantitatively measures the fraction of the Fourier power of a signal that is due to nonlinear (specifically quadratic) interaction.

Thus, the bicoherence sheds light on nonlinear coupling; whereas the bispectrum yields information regarding the energy flow due to nonlinear coupling, given the wave amplitudes in the system; although it does not by itself tell use the direction of the energy flow. It is therefore useful to compute them both when diagnosing possible nonlinear wave physics. A large value of bicoherence (close to unity) may reveal waves which have significant coupling, but do not drive additional waves in practice due to their relatively low amplitudes. This becomes apparent if one supplements the information given by the bicoherence with the bispectrum, because the latter also incorporates information about relative wave amplitudes. Conversely, plotting the bispectrum alone does not necessarily yield information about the intrinsic strength of coupling between waves.

In a 1D3V PIC simulation, in which information on both the frequency and one dimensional wavenumber spectrum is readily available, one must decide how best to calculate and display the results of bispectral analysis. In section 4.3 of chapter 4 and section 6.3 of chapter 6, the auto-bispectrum/auto-bioherence are calculated. The signal is first Fourier transformed in space and then the bispectral analysis is performed by successive Fourier transforms in time. For purely perpendicular propagation and restricting the frequencies of interest to be below the lower hybrid frequency, there is approximately a one-to-one mapping between frequency and wavenumber. This allows one to select the frequency for which the spectral power of a given wavenumber is at its maximum, which, for ICE, lies along the fast Alfvén branch. This is repeated for all wavenumbers, allowing one to calculate the bicoherence using successive Fourier transforms in time, while displaying the results of the calculation as a function of wavenumber [Irvine, 2018]. In general, PIC simulations, and certainly PIC simulations of the MCI, have much better spectral resolution in the wavenumber domain than they do in the frequency domain. The resulting bispectrum/bicoherence plot is significantly less coarse than if we had plotted it as a function of frequency.

A different approach is adopted in section 5.2 of chapter 5, in which the cross-bicoherence between two signals is calculated. First, a small segment of the total time series is selected and its average calculated. The result is a 1D signal containing the quantity of interest at a given time, as a function of position. The successive Fourier transforms are then performed in the spatial domain only. The

result is again a plot of the bispectrum/bicoherence as a function of wavenumber, but this is more coarse than the plot resulting from the other method of computing the bispectrum/bicoherence described in the preceding paragraph. The advantage of this approach, is that it is easier for one to diagnose at what time a nonlinear feature of interest “switches on”, which of course could be achieved by looking at the dispersion relation alone, but using this method one can ascertain *what* switched it on.

Chapter 3

Simulations of the MCI - Resolution testing and Parameter scans

Issues arising from the discrete nature of finite difference schemes and PIC codes were highlighted in Sec. 2.1.4. This chapter is dedicated to discussing some of these issues in more detail, aided by the results of multiple EPOCH simulations. We begin by studying how the size of the simulation domain can affect simulation results, followed by an exercise in pushing the cell size to its limits, seeing how large it can be made before the electron self-heating produces non-physical results.

The problem chosen for this investigation is similar to that of Refs. [Cook et al., 2013] and [Carbajal et al., 2014], which detail simulations of the MCI with parameters characteristic of the JET outer midplane edge using 1D3V PIC and hybrid-PIC codes respectively. The simulations consist of three plasma species, majority thermal deuterons and electrons, which are represented by Maxwellian velocity distributions, and minority suprathermal 3.5MeV alpha particles whose velocity distribution is that of the ring-beam found in Eq. 1.25. The following parameters are constant for all simulations in this chapter: electron number density $n_e = 1.0 \times 10^{19} \text{m}^{-3}$, equal deuteron and electron temperatures $T_D = T_e = 1 \text{keV}$, magnetic field $B = 2.1 \text{T}$ oriented at 88° to the simulation domain, x. The minority alpha distribution function is oriented along the magnetic field, and the concentration of alpha particles is $\zeta = n_\alpha/n_D = 10^{-3}$. The simulation domain length is varied in the simulations in this chapter, which leads to a variation in wavenumber resolution. In all cases, 200 output files are “dumped” every alpha gyro-period, leading to a maximum resolvable frequency of $100\omega_{c\alpha}$.

3.1 Domain size

This section focusses on the impact of varying the length of the simulation domain. For a grid of length L , the discretisation combined with the imposed periodic boundary conditions means only wavenumbers with $|\mathbf{k}| = n\frac{2\pi}{L}$ are represented in the simulation; where n is an integer corresponding the cell number. As the simulations presented throughout this thesis are in one spatial dimension only, henceforth the vector notation will be dropped. The Nyquist criterion restricts the maximum resolvable wavenumber in the simulation to be $n\pi/L$. So, for a simulation with n grid cells, there are $n/2$ positive wavenumbers with spacing $\Delta k = 2\pi/L$ and $n/2$ negative wavenumbers with the same spacing. This spacing Δk must be small enough to resolve the modes of interest (see Sec. 2.1.5), meaning for a periodic simulation, a large L is often required. To study the MCI, we must also be sure to resolve individual cyclotron harmonics. Let us consider an example based on the dispersion plots shown in Refs. [Cook et al., 2013] and [Carbajal et al., 2014]. For modes at low frequencies in the range, say, $\omega \lesssim 6\omega_{c\alpha}$, the fast Alfvén branch, along which the cyclotron resonances occur, has an approximately linear dispersion relation, i.e. $\omega \sim V_A k$. For these plasma parameters, we have $V_A/\omega_{c\alpha} \sim 0.1m$, meaning for every Δk we move along the wavevector axis, we move $\sim 0.1\Delta k\omega_{c\alpha}$ up the frequency axis along the fast Alfvén branch. Before even considering temporal resolution, we already need to be careful to choose our L to be large enough, and hence Δk to be small enough, so as not to “skip over” cyclotron harmonics. It is clear that we can choose plasma parameters such that the ratio $V_A/\omega_{c\alpha}$ is minimised to improve our result.

Figures 3.1, 3.2, and 3.3 show how the energy, dispersion relation, and power spectra respectively of three different simulations of the MCI under JET like conditions vary with L . The number of particles per cell was kept constant at 100, and all simulations use higher order spline shape functions. The size of the domain was altered and the number of grid cells was altered accordingly such that the width of the cell Δx remained $\sim 0.95\lambda_D \sim 0.95 \times 7.44 \times 10^{-5}m$. The sizes of the domains given as multiples of the alpha particle Larmor radius $r_L = 0.129m$ are displayed at the top of each subplot, and for each figure, the subplot axis span the same range of values. The energy transfer in each simulation, from the minority energetic ions to the thermal particles and electromagnetic fields, is near identical, and we see the oscillation of energy between the majority ions (red) and the z-component of the magnetic field (green) becomes more coherent as L increases. The dispersion plots all reveal two main clusters of excited modes along the fast Alfvén branch, one at

intermediate (k, ω) between $l = 9$ and $l = 16$ (where $\omega = l\omega_{c\alpha}$), and one at high (k, ω) between $l = 23$ and $l = 26$. The difference between the left most and right most plots is striking. The former has $\Delta k \sim 0.72 [\omega_{c\alpha}/V_A]$, leading to a pixelated looking dispersion plot. In the middle plot we can just about distinguish between neighbouring cyclotron harmonics, and the right plot is very well resolved. Figure 3.3 shows the Fourier power in the z-component of the magnetic field as a function of frequency. This type of plot is arguably the most revealing, and the most relevant for comparison with experiment. Here, only the Fourier power in intermediate (k, ω) modes is shown. In the left most panel, we see a double peak at around $l = 10$ and $l = 10.5$. This is because the mode that should be excited does not exist in the simulation, so instead the power is partitioned between the two closest modes. In the middle panel, we see this double peak merging into one, and in the panel on the far right, we see sharply defined peak, at a location somewhere inbetween $l = 10$ and $l = 10.5$ - the mode the simulation wanted to access now exists. This peak is by no means insignificant, it has around one order of magnitude more Fourier power than its neighbour at $l = 9$, and a large two orders of magnitude more Fourier power than its counterpart in the under resolved simulation displayed in the left panel. We can see a similar effect for modes at around $l = 14$ and $l = 15$. The right most plot with largest L is much smoother than the other two, owing to an increased number of data points between neighbouring modes in wavenumber space. Crucially, increasing our frequency resolution would not have helped us improve the result shown in the left panel of Fig. 3.3.




Figure 3.1: Change in energy density of particles and electromagnetic fields as a function of time. Time is plotted in units of the alpha particle gyro period. The initial energy densities of the E_x , B_y , and ΔB_z field component are all zero; while the initial energy densities of the electrons, thermal deuterons, and NBI deuterons are $\sim 2.4 \times 10^3 \text{ Jm}^{-3}$, $\sim 2.4 \times 10^3 \text{ Jm}^{-3}$, and $\sim 5.6 \times 10^3 \text{ Jm}^{-3}$ respectively.




Figure 3.2: Frequency-wavenumber space of the MCI. Shading indicates the spectral power in B_z field component and is plotted on a \log_{10} scale.

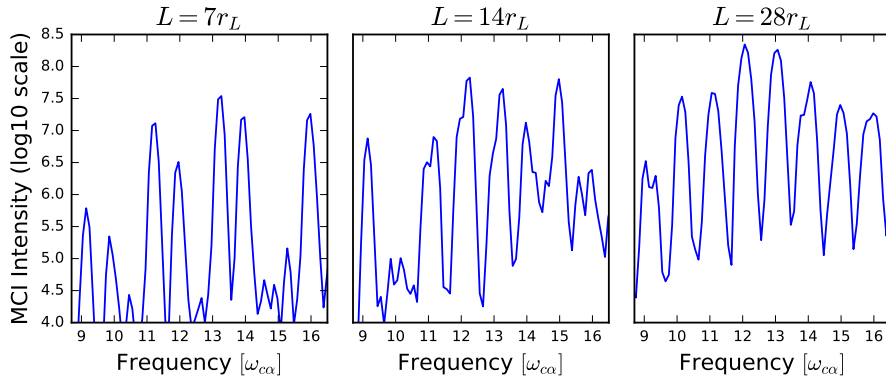


Figure 3.3: Spectral intensity of the B_z field component of the MCI simulations. Power is obtained by integrating between $0[\omega_{c\alpha}/V_A] < k < 35[\omega_{c\alpha}/V_A]$.

3.2 Cell size

We have just seen the importance of having a large domain size L , and hence a small Δk . In early PIC codes this could often pose a problem, because the cell size was essentially fixed as it was required to be some fraction of the Debye length λ_D so as not to produce non-physical self-heating. This means that if one doubles L , one must also double the number of grid cells n_x , which approximately doubles the computational expense. In more modern PIC codes, it has been shown that the electron self-heating can be ameliorated through the use of higher order shape functions and/or current smoothing [Arber et al., 2015; Buneman, 1993]. In this section we will examine the effect of increasing the cell size to multiples of the Debye length, in conjunction with two shape functions of different order.

In the following figures, panels in the left column, labelled a), b) and c), correspond to simulations using triangular shape functions, the default in EPOCH. Plots in the right column, labelled d), e) and f), correspond to simulations using higher order splines, which in EPOCH, can be turned on in the Makefile. Each row plots the results of simulations with the same cell size Δx . In row one, panels a) and d), $\Delta x \sim \lambda_D$. In row two, panels b) and e), $\Delta x \sim 3\lambda_D$, and in row three, panels c) and f), $\Delta x \sim 5\lambda_D$. In all simulations the simulation domain length is constant, and is set to $L \sim 28r_L$, which also fixes Δk . The number of cells n_x thus reduces as Δx increases. The number of particles per cell is constant at 100. This means panel d) in the following plots is identical to the right most plots in the preceding section.

First, we turn our attention to the transfer of energy, shown in Fig. 3.4. In all panels, the black line, corresponding to the change in electron energy density, increases with time. We might expect some increase in this energy, as there are electrostatic waves in the simulation with a finite parallel wavenumber, and hence the possibility of Landau damping [Cook et al., 2010; Cook, 2011]. However, we expect the evolution of this change in energy to broadly follow the evolution of the electrostatic E_x component of the field energy. In panel d) this is approximately the case, and the electron energy seems to saturate by the end of the simulation, but in all other panels, which correspond to simulations that either do not resolve the Debye length or use triangular shape functions, the electron energy has an initial increase. In panels e) and f), it looks as though the electron energy is approaching saturation by the end of the simulation, and in panels b) and c), the change in electron energy density is so rapid and large that it is not shown. This kind of behaviour can be expected in simulations using triangular shape functions and having such large cell sizes. When studying the MCI, electrons often have little to no bearing on the dynamics of the dominant field components and the ions, hence the use of hybrid-PIC codes [Carbajal et al., 2014]. So perhaps we can forgive the increase in electron energy provided the rest of the field and particle energy components evolve as we expect. However, it is clear from panels b) and c) that this is not the case. In both panels, the final energy of the B_z field component is lower than in panel a) (although its saturation energy is roughly the same), and the majority deuteron energy is much higher than in panel a). In panel c), the majority deuteron energy is still growing. Plots d), e), and f) show that using higher order shape functions is of tremendous benefit as far as the electron self-heating is concerned. The ion physics remains largely unchanged in panel f), despite the cell size being an enormous five Debye lengths.

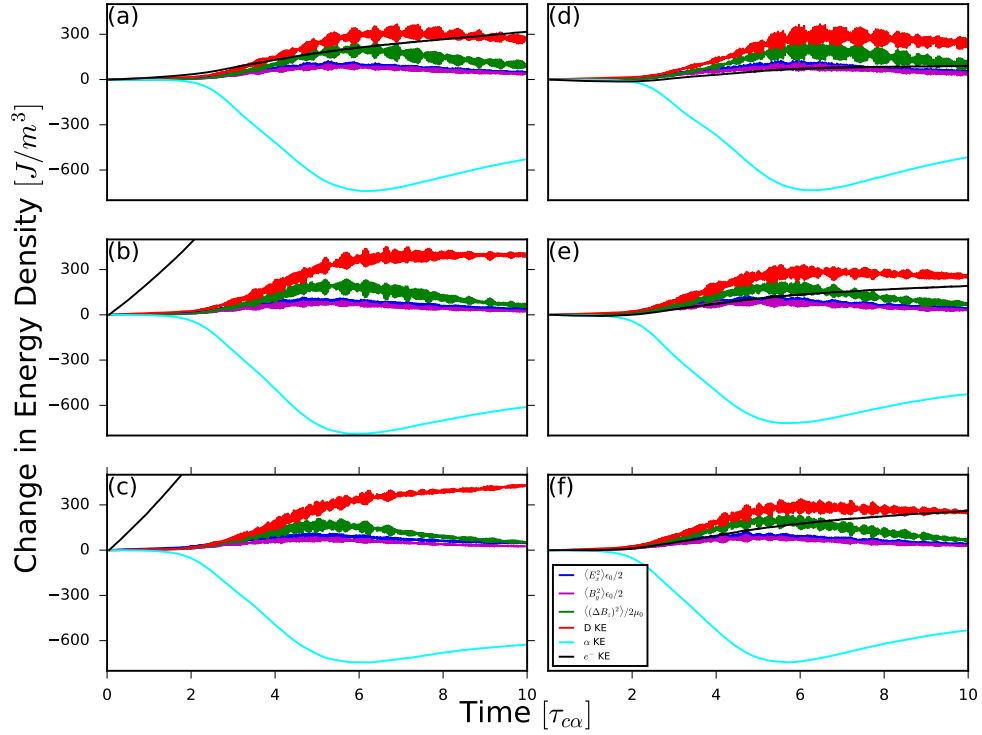


Figure 3.4: As Fig. 3.1. Panels a), b) and c) correspond to simulations using triangular simulations, panels d), e) and f) correspond to simulations using higher order spline shape functions. Row one, panels a) and d): $\Delta x \sim \lambda_D$. Row two, panels b) and e): $\Delta x \sim 3\lambda_D$. Row three, panels c) and f): $\Delta x \sim 5\lambda_D$.

We will now examine the power spectrum as a function of frequency, which is shown in Fig. 3.5 for a large frequency range. In the following discussion, it is important to bear in mind that Δk is constant across all six simulations, they have access to the same modes. The overall mode structure is the same across all simulations, but there are a few notable differences. The high frequency modes become less distinct as the cell size is increased. Looking at modes in the range $l = 9$ to $l = 16$, the structure is surprisingly robust as the cell size increases, with some changes in the relative heights of the peaks. The most striking difference is the structure of the modes with $l \lesssim 6$. These are modes that, in general, arise from the nonlinear stage of the MCI due to nonlinear interactions between modes with higher frequencies [Carbajal et al., 2014; Carbajal, 2015]. The structure is clear in panel a), becoming less so as the cell size is increased in panels b) and c). These low modes are also visible in panel d), e) and f), showing that the use of higher order shape functions helps preserve the mode structure as the cell size is increased.

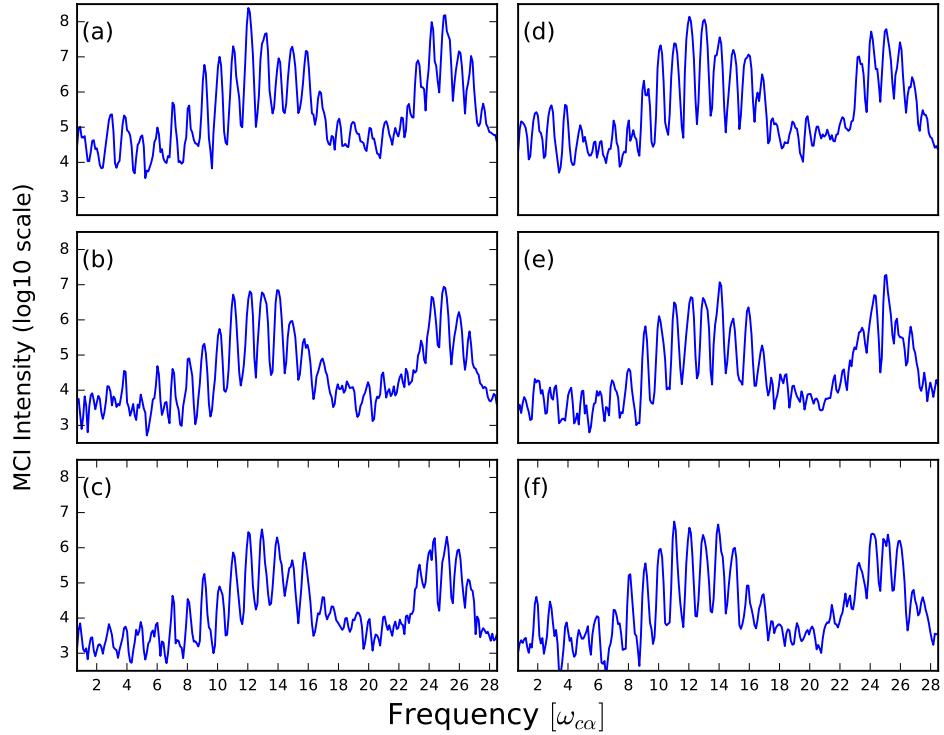


Figure 3.5: Spectral intensity of the B_z field component of the MCI simulations. Subplots are the same as in Fig. 3.4.

It is natural to consider the energy conservation properties of the code as we increase the cell size, and hence knowingly increase the amount of electron self-heating. The percentage change in the total energy of the simulations is shown in Fig. 3.6. Here, only changes in energy associated with the fluctuating part of the applied magnetic field are included in the calculation. When the applied magnetic field energy is included, the change is always a fraction of a percent. We can clearly see that simulations using triangular shape functions suffer greatly from a lack of energy conservation as the cell size further exceeds the Debye length. Panels a) and d) look very similar, the latter showing a slight improvement in energy conservation. Panels b) and c) show the total simulation energy is increasing approximately linearly. Panels e) and f) look almost identical to panels a) and d), making it abundantly clear that at least for this problem, higher order shape functions must be used if one is to have any hope of replicating the physics “cheaply” by increasing the cell size.




Figure 3.6: Percentage change in total simulation energy as a function of time. Time is plotted in units of the alpha particle gyro period. The applied magnetic field $|\mathbf{B}| = 2.1\text{T}$ has been excluded from the calculation. Subplots are the same as in Figs. 3.4 and 3.5.

3.3 Summary

In this chapter we have investigated the importance of the domain and cell sizes for an example 1D3V PIC simulation of the MCI. We found that the length of the simulation domain must be large enough to have a sufficiently small Δk in order to resolve individual cyclotron harmonics along the fast Alfvén branch. Practically, this poses a problem, as halving Δk requires us to double the domain length L , which, if we want to make sure the Debye length is resolved, requires us to double the number of grid cells n_x , thus approximately doubling the computational expense. A “cheap” way to halve Δk would be to simply double the cell size, and risk the resulting electron self-heating negatively impacting our simulation. As in Ref. [Arber et al., 2015], we have found that using higher order particle shape functions greatly reduces the electron self-heating effect in this simulation, and for this particular problem, did not have a huge impact on the ICE spectra. From this we can tentatively

conclude that it is possible to increase the cell size beyond the Debye length as long as higher order shape functions are used, at least for this particular problem. As with any computational physics problem, resolution tests such as the ones in this section must be carried out. It should also be noted that in a PIC code, due to the CFL condition, the timestep is directly proportional to the cell size, which means that doubling the cell size doubles the timestep. So if L is already fixed, halving the number of cells in theory speeds up the simulation by a factor four (if we crudely assume perfect scaling), because there are half as many particles to push in the simulation, and we can push them twice as fast. Such PIC simulations, even those in only one spatial dimension, necessarily run on multiple processors using MPI. In general, PIC simulations with larger domains scale better across multiple processes (the MPI parallelisation is achieved via domain decomposition), because what is gained through the efficient division of work outweighs the communication overhead. So practically, if one has a large number of processors at their disposal, there is often no benefit to be gained by running a “cheap” simulation, that is, one in which the wavenumber spectrum is under resolved, because such a simulation may not be able to fully exploit the processors given to it.

Chapter 4

Simulations of ICE in KSTAR

In this chapter we shall examine ICE observed in KSTAR deuterium plasmas during the third terminal stage of the ELM cycle described in Sec. 1.3. Three different categories of ICE have been observed in KSTAR, all of them at the outer midplane edge: 1) proton chirping ICE, 2) deuteron chirping ICE, and 3) deuteron steady-state ICE. In this chapter we discuss simulations pertaining to proton chirping ICE and deuteron steady-state ICE, leaving the deuteron chirping ICE as an avenue for future work.

4.1 Fusion born proton frequency chirping

4.1.1 ELMs and ICE in KSTAR

Figure 4.1 shows an example of downward ICE chirping from a KSTAR plasma that has with toroidal magnetic field at the magnetic axis $B_0 \simeq 1.99$ T, plasma current $I_p \simeq 600$ kA, and total stored energy $W \sim 380$ kJ. The upper panel displays the measured RF burst spectra as a function of time while the lower panel displays the Fourier transform of this signal. In the lower panel of Fig. 4.1, and all subsequent figures of this type, horizontal white lines denote successive proton cyclotron harmonics evaluated for the magnetic field strength at the outer mid-plane edge of the plasma, while $t = 0$ refers to the centre of a $200\mu s$ segment of RF data. This data is obtained when the RF signal amplitude exceeds a threshold voltage during KSTAR pulse operation, with the acquisition times corresponding roughly to a spike in the D_α signal [Thatipamula et al., 2016]. For further details of the fast RF spectrometer system used on KSTAR see for example Fig. 1 of [Thatipamula et al., 2016].

The frequency chirping observed in KSTAR is often, as in Fig. 4.1, in discrete steps coinciding with the local proton cyclotron frequency. The only energetic




Figure 4.1: Temporal evolution of ICE amplitude (upper plot) and spectrum (lower plot) during an ELM crash in KSTAR plasma 11513. Time is measured relative to the moment chirping bursts are observed to begin during the ELM crash ($\sim 100 \mu s$ after the start of the crash). The horizontal dashed lines in the spectrogram indicate proton cyclotron harmonics.

protons in KSTAR plasmas are those produced in deuterium-deuterium (D-D) fusion reactions, hence it is likely that ICE at spectral peaks separated by proton cyclotron harmonics is driven by fusion-born protons. If this ICE is driven by the MCI of confined fusion-born protons with spatially localised population inversion, it is necessary to identify a candidate population. The KSTAR experiment is not built on a scale sufficiently large to confine the majority of the energetic ions that are born in fusion reactions within pure deuterium plasmas: 3.0 MeV protons, 0.82 MeV He-3 nuclei, and 1.0 MeV tritons. To drive ICE in KSTAR via the MCI, there must nonetheless exist a collectively unstable subset of fusion-born protons whose orbits are confined within the plasma and traverse the excitation region in the outer midplane edge. The distribution in velocity space of this population must differ qualitatively from that driving ICE in other MCF plasmas such as JET and TFTR, because the fusion-born ions on the marginally trapped orbits that give rise to ICE in these two larger

experiments are promptly lost from KSTAR during their first drift excursion.

4.1.2 Single particle orbits

We first calculate orbits of 3.0 MeV protons in equilibrium magnetic fields resembling those of a KSTAR plasma with major radius $R_0 = 1.8$ m, toroidal magnetic field $B_0 = 2.27$ T, and plasma current $I_p = 611$ kA. Our orbit calculations show that almost all centrally-born fusion protons are lost promptly from the plasma on their first drift orbit. However, a small fraction of these protons are born onto deeply passing orbits which remain confined at the plasma edge to which they are carried by their radial drift excursion. Figure 4.19 shows example orbits of 3.0 MeV protons born in the midplane at initial major radii $R(0)$ equal to (a) 1.85 m and (b) 1.90 m, and with initial velocity vectors slightly offset from the co-current toroidal direction. These orbits pass through the outer midplane plasma edge, and could give rise there to a localised population in velocity space. We show below that such a distribution can radiate collectively through the MCI and is thus capable of driving ICE at proton cyclotron harmonics characteristic of the outer midplane. We conjecture that when the ELM crash starts, confinement of all energetic ions at the edge is lost; the edge is then rapidly re-populated on a drift orbit timescale by this subset of the energetic fusion-born protons. This leads to the sharp local population inversion that drives ICE. We note that an instance of positive correlation between ELMs and ICE was also seen on JET, see Fig. 9 of [Cottrell et al., 1993].

The spatial location from which ICE is emitted is inferred from the observed frequency interval between successive spectral peak features, which corresponds to the local proton cyclotron frequency. This fixed proton cyclotron frequency implies that the radial location of the emitting region does not change. The full spectral evolution takes place on a microsecond timescale which is very short compared to any characteristic timescale of evolution of the energetic ion distribution in velocity space; thus we do not attempt to model any changes in the energetic ion distribution function. The KSTAR magnetic field strength uniquely defines a radial location for the protons driving the ICE, to high accuracy, at the outer mid-plane.

4.1.3 Self consistent simulations of the MCI

To simulate the excitation of ICE by fusion protons in KSTAR we use the EPOCH PIC code as described in chapter 2. As the simulations are set up in planar geometry, they do not incorporate realistic toroidal geometry and the associated compressional Alfvén eigenmode structure [Fülöp et al., 2000; Gorelenkov and Cheng,




Figure 4.2: Poloidal projection of 3.0 MeV fusion proton orbits in the model KSTAR equilibrium with initial velocity vectors slightly offset from the co-current toroidal direction. The blue boxes in both the left and right panels designate the emitting region for the ICE and are centred on the radial location that corresponds to the proton cyclotron frequency whose harmonics are excited in the chirping ICE spectrum. The horizontal extent of these boxes corresponds to the local Larmor radius of cyclotron gyration of an emitting proton.

1995; Gorelenkov, 2016; Gorelenkov et al., 2014; Kolesnichenko et al., 1998, 2000; Smith et al., 2003; Hellsten and Laxaback, 2003; Hellsten et al., 2006; Smith and Verwichte, 2009]. Nevertheless this approach has been shown to be successful in capturing most of the underlying physics [Cook et al., 2013; Carbajal et al., 2014, 2017], and aligns with the original slab-geometry analytical theory of the MCI described in Sec. 1.4, against which it has been benchmarked, see Fig. 1 of [Cook et al., 2013] and Fig. 1 of [Carbajal et al., 2014]. The simulations capture the rapid self-consistent time-evolution of the ion velocity distributions under the MCI, but do not attempt also to capture the much slower (by orders of magnitude) effects of collisional slowing-down.

The deeply passing subset of the fusion-born proton population in the KSTAR plasma edge discussed above has a speed perpendicular to the magnetic field ($v_{\perp 0}$) much smaller than the magnitude of their speed at birth. It is therefore justifiable to represent this by the approximate model of a delta-function ring distribution. Previous analytical and first principle, fully self-consistent simulation studies of the MCI

showed that strong driving of the MCI is enabled if the perpendicular component of the energetic ion velocity is at least comparable to the local Alfvén velocity. In all our simulations, $v_{\perp 0}$ corresponds to an energy $150 \text{ keV} \simeq 5\%$ of the birth energy consistent with the orbit calculations shown in Fig. 4.19. This is comparable to the local Alfvén speed c_A , hence high enough to drive the MCI. The large parallel velocities v_{\parallel} of these passing fusion protons are not represented in the simulations since the parallel dynamics of the fast ions play no role in perpendicular wave propagation. The ratio of proton number density to deuteron number density is 10^{-3} and the total simulation duration is 10 proton cyclotron periods, carrying the MCI into its saturated nonlinear regime. The spatial (x) direction is orthogonal to the uniform magnetic field $\mathbf{B} = B_{\text{cyc}} \mathbf{e}_z$, so that the propagation direction of waves excited in the system is perpendicular to \mathbf{B} . B_{cyc} denotes the edge magnetic field inferred from the spacing between successive proton cyclotron harmonics observed in the experimental spectrograms. The bulk plasma comprises electrons and deuterons with initial temperature 1 keV.

The distribution of collectively radiated energy between different cyclotron harmonics depends on the character of the instability (the MCI) driving the ICE. A key dimensionless parameter in the analytical theory of the MCI is the ratio of energetic ion perpendicular velocity to the Alfvén speed. Changing the density changes the Alfvén speed, which in turn changes the numerical value of this ratio. This ratio determines the threshold for instability which can be inferred from both linear theory and fully nonlinear self-consistent PIC simulations. These nonlinear simulation results have many good points of contact [Cook et al., 2013; Carbajal et al., 2014, 2017] with the linear instability theory of the MCI, see in particular, Figs. 1-3 in Ref. [Carbajal et al., 2017]. This means the character of the MCI depends strongly on the local density so that if the local density changes, the set of cyclotron harmonics which are excited by the MCI also changes. Thus the spectral character of the ICE at any time reflects the density-dependence of the MCI.

To enable comparison with observations of ICE chirping during the ELM crash in KSTAR plasmas, we run multiple simulations to determine the density-dependence of the MCI in the fully nonlinear saturated regime. These simulations are carried out with initial electron number densities in the range $0.2 \times 10^{19} \text{ m}^{-3}$ to $2.5 \times 10^{19} \text{ m}^{-3}$. This range reflects Thomson scattering measurements in the edge pedestal (see the last paragraph of section 3 in Ref. [Thatipamula et al., 2016]).

4.1.4 ICE intensity as a function of frequency and electron number density

Figures 4.3 and 4.4 show the spectra of saturated ICE intensity obtained from EPOCH simulations with successive values of electron number density (lower panels), along with the corresponding experimentally-measured spectrograms for downward chirping ICE during ELM crashes in plasmas with $B_0 = 1.7\text{T}$ and $B_0 = 1.99\text{T}$ (upper panels). The lower left and lower right panels correspond to the high and low frequency ranges respectively. Electron number density decreases from left to right, and each vertical strip corresponds to an independent MCI simulation at the density shown. The arrows labelled (a)-(d) denote a mapping between experimentally observed and simulation proton cyclotron harmonics. The boundary of evanescence in the cold plasma limit, defined by the lower hybrid frequency ω_{LH} , is clearly visible. As the electron number density, and subsequently ω_{LH} decrease, the number of modes available for excitation also reduces. For $n_e \lesssim 1.1 \times 10^{19}\text{m}^{-3}$, the spectrum is dominated by a single cyclotron harmonic, with the harmonic number falling monotonically as n_e decreases. It is apparent that for both values of central magnetic field B_0 shown in Figs. 4.3 and 4.4, the trend in ICE spectral power as a function of frequency and n_e is the same.

The variation of simulated ICE intensity with electron number density resembles the experimentally-observed variation of ICE intensity with time. We address this in greater detail in the next section, meanwhile we note that the striking agreement between the highest and lowest intensity regions in the upper and lower panels of both Figs. 4.3 and 4.4 reinforces the choice of simulation electron number density values which reflect Thompson scattering measurements. There is a “missing harmonic” in the lower panels of Figs 4.3 and 4.4 at $f \simeq 323\text{MHz}$ and $f \simeq 375\text{MHz}$ respectively. This could be due to periodic boundary conditions, as well as the limitations of a 1D3V model.

4.1.5 Mapping between computational results and KSTAR observations

Our simulations suggest that the downward chirping of ICE observed during the KSTAR ELM crashes detailed above is likely to be a direct consequence of a rapid local decline in electron number density associated with an ELM filament. An important corollary of this result is that measurements of the ICE spectra can be used to infer the locally declining electron number density as plasma is transported out of the ICE-emitting region by one of the multiple filament bursts that are associated




Figure 4.3: Top panel: experimentally-measured fast RF burst spectrogram from KSTAR plasma 11462 with $B_0 = 1.7T$ and average electron number density before the ELM crash $\langle n_e \rangle = 2.5 \times 10^{19} m^{-3}$. Downward step-wise frequency chirping with proton cyclotron frequency $f_{cp} \sim 21.5 MHz$ is apparent. Lower panels: frequency versus electron number density plots for the nonlinear stage of MCI simulations where $B_z = B_{cyc} \approx 1.41T$ has been inferred from the data in the top panel. Shading indicates the \log_{10} of the spectral power in the fluctuating part of the B_z field component of each simulation.

with the ELM crash.

In Fig. 4.4 (upper panel), a given ICE spectral feature at a particular cyclotron harmonic in KSTAR can be seen to arise at a time t_{start} and end at a time t_{finish} . For this downward chirping case, a neighbouring spectral feature at a lower cyclotron harmonic arises at a slightly later t_{start} and persists until a slightly later t_{finish} . In the lower panels of Fig. 4.4, a succession of simulated ICE spectra are shown. Each vertical strip represents an ICE simulation run at a different, neighbouring, density. These simulated spectra are arranged in sequence with number density declining from left to right. A particular cyclotron harmonic spectral feature (e.g. in the lower right of Fig. 4.4) is excited by simulations in a relatively narrow range of densities between n_{upper} and n_{lower} . Remarkably, the pattern of spectral

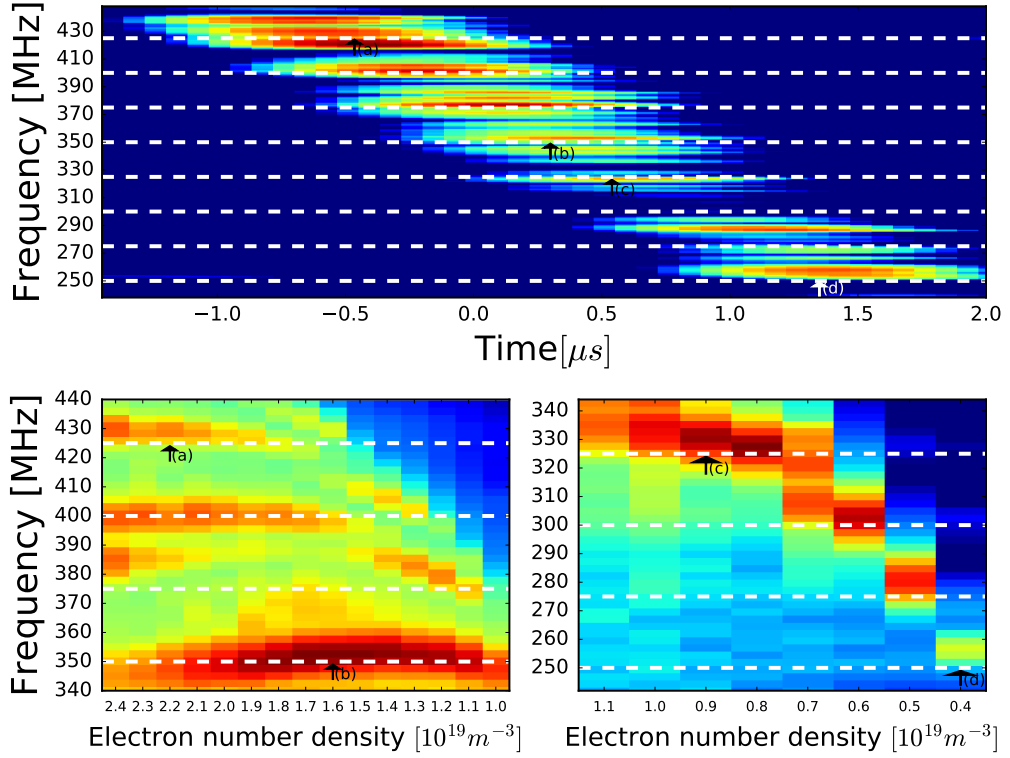


Figure 4.4: Top panel: experimentally-measured fast RF burst spectrogram from KSTAR plasma 11513 with $B_0 = 1.99\text{T}$ and average electron number density before the ELM crash $\langle n_e \rangle = 2.6 \times 10^{19}\text{m}^{-3}$. Downward step-wise frequency chirping with $f_{cp} \sim 25\text{MHz}$ is apparent. Lower panels: as the lower panel of Fig. 4.3 but with $B_z = B_{cyc} \approx 1.64\text{T}$.

features in the sequence of simulated spectra, depending on number density n , has much in common with the pattern of spectral features in the experimental ICE spectra, depending on time. Hence we can construct a simple mapping from time to density by exploiting these similarities. Essentially, for a given spectral feature, we identify n_{upper} with the density at t_{start} , and n_{lower} with the density at t_{finish} .

The foregoing describes the overall approach. It yields, we argue, sub-microsecond time resolution for the ELM filament burst. This is shown in Fig. 4.5, the left panel corresponding to Fig. 4.3, and the right panel corresponding to Fig. 4.4. Our sequence of simulations at different electron number densities, together with the chirping ICE observations, reflect a phase of the ELM crash during which individual, toroidally and poloidally localised filament bursts lasting $< 10\mu\text{s}$ (Fig. 4, Ref. [Thatipamula et al., 2016]) are known to be present. In greater detail, the procedure for obtaining Fig. 4.5 (left) from Fig. 4.3 is as follows:

1. Corresponding experimental and computational spectral features are identified, e.g. the feature denoted by arrow (a) in the upper panel of Fig. 4.3 is compared with the feature denoted by (a) in the lower left panel of Fig. 4.3.
2. The most intense regions of the experimental and simulation spectral features are identified and quantified with respect to time and density, respectively.
3. If a spectral feature arises intensely in “N” simulations that have different (neighbouring) values of electron number density, the corresponding experimental feature is then divided into N time points.
4. One of these N time points is matched with its corresponding density point, and a coordinate is plotted in Fig. 4.5. The vertical error bars in this plot reflect the finite steps in density between simulations, and as such are always $0.1 \times 10^{19} \text{m}^{-3}$. The width of the horizontal error bars is determined by the difference in the successive time-intervals outlined in iii).
5. In the case where we only have one density point corresponding to one spectrally dense feature, e.g. arrows (d) in Fig. 4.3, the time is taken to be the centre of the spectrally dense region in the experimental spectrogram. The width of the temporal error bars is then half the width of the spectrally dense region in the upper, experimental plot. Hence the larger error bars at low densities in Fig. 4.5.

A slightly more detailed explanation, with accompanying plots, can be found in Appendix B. This yields an available time resolution of the density decrease of order $\sim 0.1 \mu\text{s}$. Fig. 4.5 shows that the density decreases approximately linearly on a timescale of about $2.5 \mu\text{s}$.

4.1.6 Upward chirping

Upward chirping is occasionally observed during ELM crashes in KSTAR. For logical consistency with the above, we should attribute this to a rare locally rising density due to the rapid motion of an ELM filament transporting additional plasma into the ICE-emitting region during the ELM crash. We test this hypothesis as follows. The time evolution of the corresponding electron number density for the case of upward chirping shown in Fig. 4.6 (left panel) has been inferred from the saturated MCI simulations at different electron number densities (Fig. 4.6 right panel), and is shown in Fig. 4.7.

In Fig. 4.6, the separation between successive spectral peaks in the ICE is $f_{cp} \sim 28 \text{MHz}$, which corresponds to a local magnetic field strength of 1.84T and




Figure 4.5: Local electron number density during an ELM crash in KSTAR inferred from downward chirping ICE measurements compared with saturated MCI simulations. Left (right): fitted and estimated electron number density and time values corresponding to KSTAR plasma 11462 (plasma 11513). The green fit suggests that local electron number density declines approximately linearly with time during the early stages of the ELM crash.




Figure 4.6: Left panel: experimentally-measured fast RF burst spectrogram from KSTAR plasma 11474 with $B_0 = 2.27T$ and $\langle n_e \rangle = 2.5 \times 10^{19}m^{-3}$. Upward step-wise frequency chirping with $f_{cp} \sim 28MHz$ is apparent. Right panel: as the lower panel of Fig. 4.3 but with $B_z = B_{cyc} \approx 1.84T$.

thus implies that the ICE source is localised to the outer midplane of this KSTAR plasma. The absolute value of each spectral peak frequency is at integer multiples of 28 MHz plus, systematically, 10 MHz. We conjecture that this systematic frequency shift may be a Doppler shift. If, as we suggest, the source plasma for the upward chirping ICE in Fig. 4.6 is a moving ELM filament, this Doppler shift would be a consequence of filament motion. Testing and exploitation of this particular aspect of the phenomenology would require information of the k_{\parallel} spectrum of the ICE, which is not at present available.




Figure 4.7: Rising electron number density during an ELM crash inferred from upward chirping ICE measurements in KSTAR plasma 11474 combined with saturated MCI simulations at different densities, see Fig. 4.6. The green fit suggests electron number density rises approximately linearly with time during the early stages of the ELM crash.

4.2 The effect of the lower hybrid frequency

In the previous section, a series of PIC simulations at successive neighbouring fixed values of plasma electron number density n_e were used to model the chirping of the primary ICE features in Fig. 4.1. The results of this section show a clear dependence on the lower hybrid frequency $\omega_{LH} = 2\pi f_{LH}$ (see Sec. 1.2), whereby the number of modes available for excitation decreases rapidly as the electron number density n_e , and hence ω_{LH} , decreases. We now address the role of ω_{LH} , and in particular the question of evanescence for cold plasma waves at frequencies above ω_{LH} . As a guide to the reader, Fig. 4.8 shows how the lower hybrid frequency varies as a function of density and magnetic field. We see that as n_e decreases, ω_{LH} decreases rapidly.

Several more simulations were run using different values of the magnetic field, in an attempt to quantify the effects of changing the magnetic field strength B_0 as well as density n_e , and the effects of ω_{LH} on the simulation output. The magnetic




Figure 4.8: Left: The lower hybrid frequency ω_{LH} as a function of electron particle density n_e for the KSTAR plasma parameters used in the PIC simulations. ω_{LH} is normalised to the proton cyclotron frequency ω_{cp} while n_e is normalised to 10^{19} . Each line denotes ω_{LH} for different values of central magnetic field B_0 spanning 1.7T at the top (blue) to 2.3T at the bottom (black) in steps of 0.1T. Right: ω_{LH} as a function of B_0 calculated for different values of n_e spanning $2.4 \times 10^{19} m^{-3}$ at the top (blue) to $0.2 \times 10^{19} m^{-3}$ at the bottom (magenta) in steps of $0.2 \times 10^{19} m^{-3}$.

field strength in these simulations was taken to be that of the toroidal magnetic field component in KSTAR using different values of B_0 , i.e $B_z = B_\varphi = B_0 \frac{R_0}{R_0+a}$ and the small poloidal contribution is neglected. As such, the following have no exact experimental counterparts to date, but the range of B_0 simulated are applicable to KSTAR operating regimes. A total of 84 new simulations are performed, varying B_0 and n_e , the former from 1.7T to 2.3T in steps of 0.1T, the latter from $2.4 \times 10^{19} m^{-3}$ to $2.0 \times 10^{18} m^{-3}$ in steps of $2.0 \times 10^{18} m^{-3}$. The results of these simulations are shown in Figs. 4.9, 4.10 and 4.11, in all panels n_e decreases from left to right. The domain length of each simulation for a given B_0 is reduced from $50r_{L,p}$ (as was the case in Sec. 4.1) to $20r_{L,p}$ where $r_{L,p}$ is the proton Larmor radius, and the number of grid cells is in turn reduced whilst still resolving the Debye length. This reduction in the number of grid cells allows for quicker, and hence many more simulations, whilst

still ensuring the particle dynamics are well resolved. With the exception of n_e and B_0 , all other plasma parameters are constant and the same as those in the previous section, and the total simulation time is again 10 proton cyclotron periods in all cases. As before, the magnetic field of each simulation is oriented entirely along the z-axis meaning we restrict the study to strictly perpendicular wave propagation along the x-axis, which is the spatial domain of our 1D3V PIC simulations. As before, each vertical strip corresponds to an independent simulation which yields the spectrum of MCI-excited waves at the value of n_e shown; and shading indicates the \log_{10} of the Fourier power in the B_z component of the simulation.

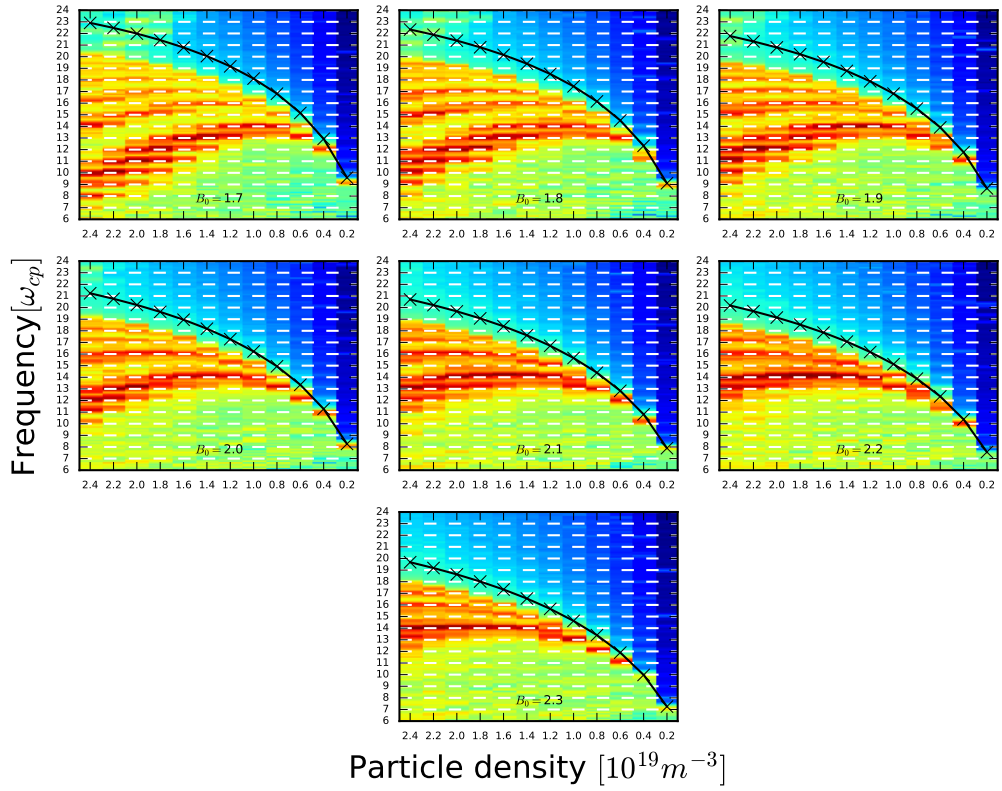


Figure 4.9: Spectral power as a function of plasma density obtained from multiple PIC simulations for several values of magnetic field strength corresponding to the nonlinear saturated phase of the MCI. Shading indicates the \log_{10} of the spectral power in the B_z field component of each simulation. All panels comprise a series of simulations, each contributing a vertical strip at a different density. The white lines denote successive proton cyclotron harmonics, while the black lines denote the value of the lower hybrid frequency (in units of ω_{cp}) for each value of density.

Figure 4.9 shows the variation of ICE intensity as a function of frequency and density for several values of B_0 which are stated in the lower centre of each panel. Frequency is in units of the proton cyclotron frequency, the white lines denote successive proton cyclotron harmonics, and the black line denote the value of the lower hybrid frequency (in units of ω_{cp}) for each value of density.

Figure 4.9 highlights several striking features, the first being that the general trend in the ICE spectral power as a function of frequency and electron particle density is the same for all values of magnetic field. It is apparent that the number of modes that can be excited for a given value of n_e reduces as the magnetic field increases, this may be due to the lower value of ω_{LH} (when in units of ω_{cp}) restricting the number of available modes for excitation. It can also be seen that the difference between the highest and lowest excited proton cyclotron harmonics reduces as the value of the magnetic field increases. In all panels, there is value of density for which only one frequency is excited at this density and all below it, this value of density increases as the magnetic field increases. We note in passing the extremely narrow range of frequencies that can be excited for $B_0 = 2.3\text{T}$, which provides further justification for attributing the results of Fig. 4.6 to upward frequency chirping. In all panels, one can see a blue region in which the spectral power falls to zero. The boundary of this region at each value of n_e lies close to the corresponding value of ω_{LH} .

An alternative way to analyse the data shown in Fig. 4.9 is to plot the variation of ICE intensity as a function of frequency and B_0 for each value of density. This is shown in Fig. 4.10, in which each panel displays the results of seven simulations each with the same density and different values of magnetic field. The black lines denotes ω_{LH} in units of the ω_{cp} for each value of B_0 . It is important to note that while the vertical frequency axis is common to all adjacent vertical strips when plotted in units of ω_{cp} , it is different when plotted in MHz, and adjacent strips can no longer share the same axis. This alternative view helps to reinforce some of our previous observations.

Figure 4.10 makes it clear that the number of frequencies available for excitation for a given density reduces as magnetic field increases. It also highlights that the number of excitable frequencies rapidly reduces as density decreases, regardless of the magnetic field. Finally, it enables us to see that when a certain value of n_e is reached, for this value and all below it, only one frequency can be excited regardless of the magnetic field present. This is approximately $8 \times 10^{18} \text{ m}^{-3}$.

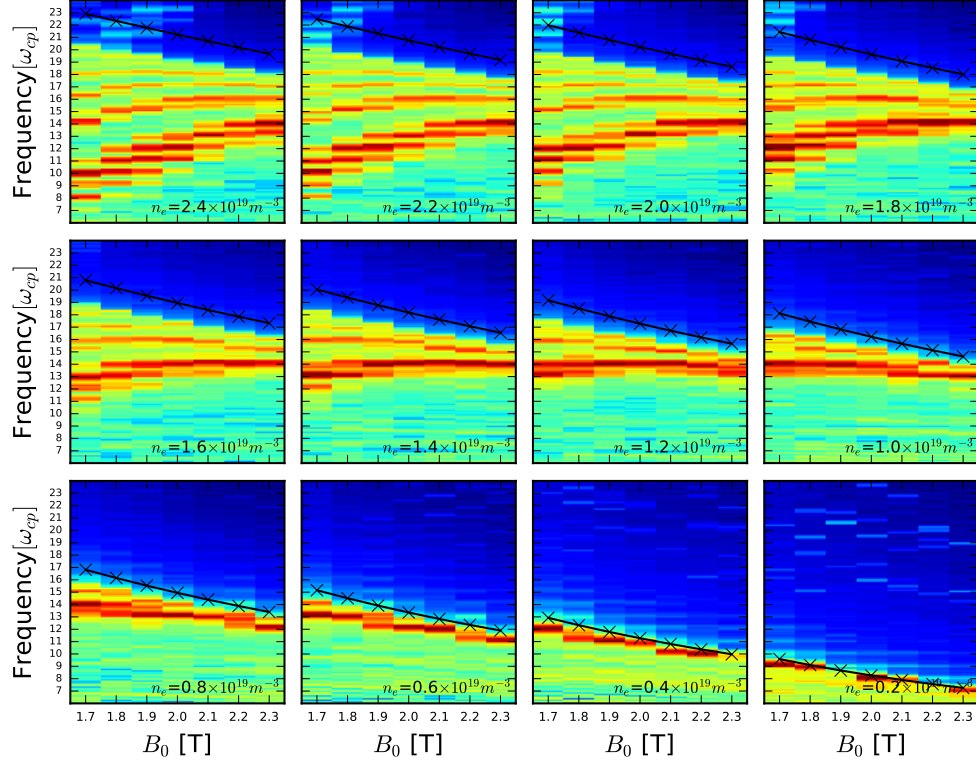


Figure 4.10: Simulated frequency vs magnetic field plots for the non-linear stage of the MCI. Frequency is in units of the proton cyclotron frequency. Each panel shows the variation of the \log_{10} of the spectral power in the B_z field component with frequency and magnetic field at constant density. The black and white lines indicate the position of the lower hybrid frequency (in units of the proton cyclotron frequency) and successive proton cyclotron harmonics respectively.

Figure 4.11 provides a further alternative with which to view the above data. Each panel displays the \log_{10} of the average spectral power in the fluctuating part of the B_z field component of each simulation harmonic as a function of n_e and B_0 . For each simulation with a given B_0 and n_e , the power in each harmonic is averaged over $\pm l/2$ where l is the harmonic number. For instance, the shading in the 6th harmonic (top left panel) for a given B_0 and n_e , corresponds to the average spectral power between $\omega/\omega_{cp} = 5.5$ and $\omega/\omega_{cp} = 6.5$, the 7th harmonic corresponds to the average power between $\omega/\omega_{cp} = 6.5$ and $\omega/\omega_{cp} = 7.5$, and so on. The black dots denote the values of B_0 and n_e for which the lower hybrid frequency is within $\pm l/2$ of the harmonic in question. This plot enables us to pinpoint the values of B_0 and n_e for which the spectral intensity of a given harmonic is highest. Combining this with future simulations which include effects such as wave dissipation could

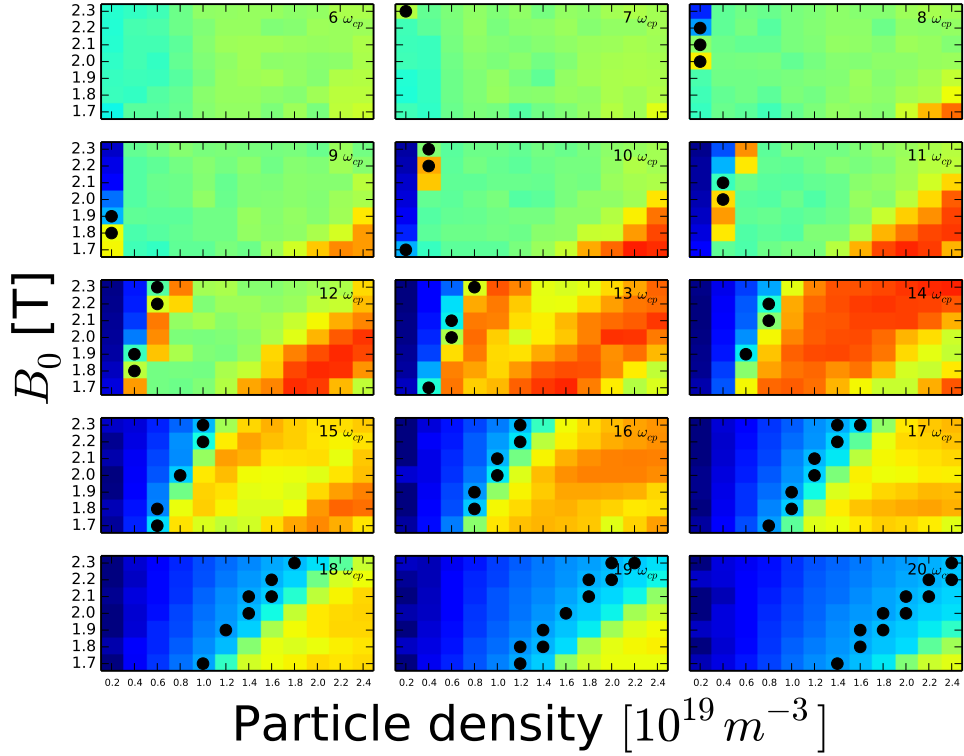


Figure 4.11: The variation of the \log_{10} of the average spectral power in the B_z field component of a given harmonic as a function of magnetic field and electron particle density. Each panel corresponds to a different simulation harmonic. Black dots denote the values of B_0 and n_e for which the lower hybrid frequency is within $\pm l/2$ of the harmonic in question, where l is the harmonic number.

yield valuable information. For instance if one is certain they are observing say, the 14th proton cyclotron harmonic at a given time, one could conjecture that B_0 must lie somewhere between 1.9T and 2.2T, and that the electron particle density lies somewhere between $6 \times 10^{18} \text{ m}^{-3}$ and $8 \times 10^{18} \text{ m}^{-3}$.

4.3 Nonlinear wave-wave interactions

The KSTAR RF signal shown in Fig. 4.1 also includes chirping features with sharply-defined spectral structure in the frequency range up to ~ 900 MHz. If our downward density chirping interpretation holds, then this is above the local lower hybrid frequency at the top of the main chirping branch, $f_{LH} \sim 529$ MHz. This “ghost” chirping feature above ~ 500 MHz is shown in Fig. 4.12, here we show that this is a real physical phenomenon, which is generated by strong nonlinear wave-wave cou-

pling between different spectral peaks within the primary chirping ICE feature below ~ 500 MHz. We demonstrate this by direct bicoherence analysis of: first, the KSTAR data files for field magnitudes; and second, analysis of the EPOCH simulation data used to make up the lower panels of Fig. 4.4. We show that the “ghost” spectral features are able to exist and grow in the higher-frequency, potentially evanescent, region because they are nonlinearly driven by coupled MCI-excited waves that lie within the lower-frequency, propagating (non-evanescent), region. The “ghost” feature thus owes its existence to both a minority suprathermal ion population - here, the confined subset of fusion-born protons that relax through the MCI in KSTAR deuterium plasmas - and on the capacity of the plasma to nonlinearly couple together the modes initially driven by these protons.




Figure 4.12: Temporal evolution of ICE amplitude during an ELM crash in KSTAR plasma 11513. Time $t = 0$ refers to the centre of a $200 \mu s$ segment of radio-frequency data. The horizontal dashed lines in the spectrogram indicate energetic proton cyclotron harmonics f_{cp} at the low field side plasma edge. In addition to the main chirping feature $\lesssim 500$ MHz $\approx 20 f_{cp}$ discussed in section 4.1, we also observe a second, faint (“ghost”), feature at frequencies above the lower hybrid frequency $f_{LH} \approx 529$ MHz $\approx 21 f_{cp}$. This additional, spectral feature is delayed in time by approximately $1 \mu s$ with respect to the main chirping feature.

4.3.1 Bicoherence analysis

We first examine the extent of nonlinear wave-wave coupling within the experimental dataset that spans the primary and “ghost” chirping ICE features in Fig.4.12. The best quantitative evidence for this coupling, and characterisation of its magnitude as a function of wave frequency, is obtained from bispectral analysis, whose properties are detailed in section 2.2. The bicoherence and bispectrum corresponding to the entire KSTAR signal shown in Fig.4.12 are plotted in the left and right panels respectively of Fig.4.13. In the bicoherence panel, shading indicates the intrinsic strength of nonlinear coupling, 1 (dark red) being completely coupled and 0 (dark blue) completely uncoupled. The shading of the bispectrum panel is displayed on a logarithmic scale. Here the averages denoted by $\langle \cdot \rangle$ in Eq.(2.19) are taken over a time window $\Delta t \sim 0.5\mu s$ within a signal which is $5\mu s$ long, corresponding to the data displayed in Fig. 4.12. This choice enables us to construct ten independent realisations. In consequence, the threshold for significance is comfortably below the observed coupling strength “b” for a wide range of relevant frequencies.

We note three distinct regions of strong intrinsic nonlinear wave coupling in the left panel of Fig. 4.13:

1. $200\text{MHz} (8f_{cp}) < f_1, f_2 < 500\text{MHz} (20f_{cp})$: Coupling between neighbouring modes within the main chirping feature shown below $f \approx 450$ MHz in Fig. 4.12. We argue that this coupling enables formation of the faint higher frequency “ghost” chirping feature that appears above $f \approx 450$ MHz in Fig. 4.12.
2. $500\text{MHz} (20f_{cp}) < f_1, f_2 < 850\text{MHz} (34f_{cp})$: Weaker coupling between the newly formed modes in the “ghost” feature above $f \approx 450$ MHz.
3. $500\text{MHz} (20f_{cp}) < f_1 < 850\text{MHz} (34f_{cp})$ and $200\text{MHz} (8f_{cp}) < f_2 < 500\text{MHz} (20f_{cp})$: Weaker coupling between the newly formed modes in the “ghost” feature above $f \approx 450$ MHz, and modes within the main chirping feature below $f \approx 450$ MHz.

We are primarily concerned with point (i), which strongly suggests the “ghost” feature is a real plasma physics phenomenon. The right panel of Fig.4.13 indicates why it is only waves in the frequency range below $f \approx 450$ MHz that can drive the




Figure 4.13: Left: Bicoherence of the observed fast RF signal displayed in Fig. 4.12, plotted as a function of frequency in MHz. The colour scale indicates intrinsic nonlinear coupling between waves with frequencies f_1 and f_2 , which takes values between 0 and 1. There is significant coupling in three distinct regions, discussed as (i) to (iii) in this section. Right: Bispectrum of the same RF signal. Bicoherence measures the intrinsic strength of nonlinear wave coupling, while the bispectrum measures actual nonlinear transfer of energy. Colour is plotted on a log scale.

observed “ghost” features: these are the waves that are not only significantly nonlinearly coupled, but also have sufficiently large amplitude. The nonlinearly driven features that could in principle arise due to the strong coupling of waves described in points (ii) and (iii) would lie below the Nyquist frequency; however, they are never observed in practice because their amplitude is lower by several orders of magnitude. We note that the auto-bispectrum and auto-bicoherence of the KSTAR RF data, that is, bispectra computed from a single time series, cannot by themselves yield information on the direction of energy transfer. To do so would require two point measurements [de Witt, 2003; de Wit et al., 1999] which at present, are not available.

Having inferred from bispectral analysis of the KSTAR data that the nonlinear wave coupling between cyclotron peaks below $f \approx 500$ MHz drives the “ghost” chirping feature, the question now arises: can the same physics be inferred from analysis of the outputs of the corresponding PIC simulations? The simulations have a propagation angle $\theta = 90^\circ$, for which, as noted above, the region $f_{LH} < f < f_2$ is evanescent. In order to explore the hypothesis that the observed waves in this region arise from nonlinear wave coupling, let us focus on the simulations which make up the lower panels of Fig. 4.4.

Figure 4.14 shows the bicoherence plots along with the corresponding spatio-

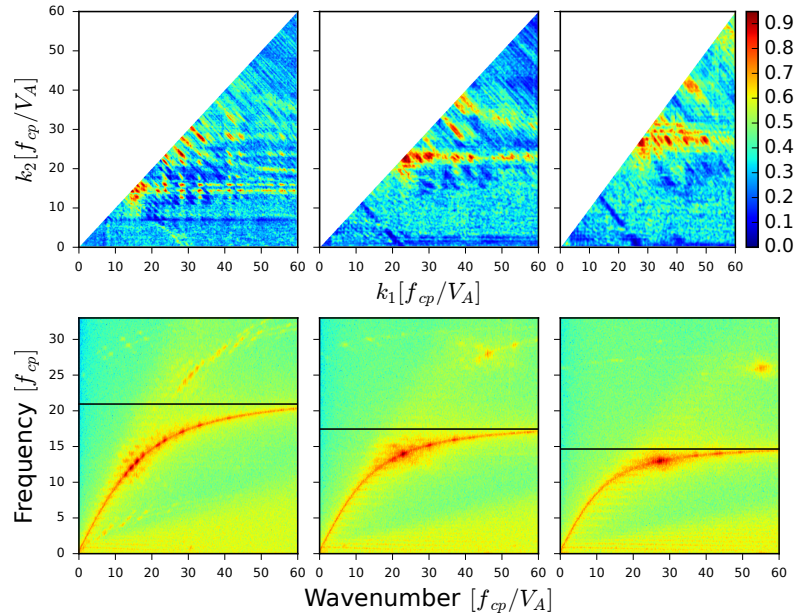


Figure 4.14: Upper panels: Normalised bicoherence (Eq. (2.19)) of the oscillatory part of the B_z field component in three PIC simulations. This is plotted as a function of wavenumber normalised to f_{cp}/V_a where V_a is the Alfvén speed. From left to right the number density n_e in the simulations is $2.4 \times 10^{19} \text{m}^{-3}$, $1.3 \times 10^{19} \text{m}^{-3}$, and $0.8 \times 10^{19} \text{m}^{-3}$. In all three plots the most significant coupling is observed between neighbouring k values near the $k_1 = k_2$ boundary. Lower panels: Corresponding spatio-temporal Fourier transform for the three simulations. The y-axis is plotted in units of f_{cp} while the x-axis is plotted in units of f_{cp}/V_a . The horizontal black line denotes the lower hybrid frequency f_{LH} for each simulation. In addition to the expected cold plasma waves below f_{LH} , spectrally intense regions above f_{LH} can be seen in the range of frequencies corresponding to the observed “ghost” feature in Fig. 4.12.

temporal Fourier transform of B_z for each of three different simulations in the lower panels of Fig. 4.4. Shading indicates the \log_{10} of the spectral density of the oscillatory part of the B_z field component. From left to right the simulations have electron number densities $n_e = 2.4 \times 10^{19} \text{m}^{-3}$, $1.3 \times 10^{19} \text{m}^{-3}$, and $0.8 \times 10^{19} \text{m}^{-3}$. In the lower panels, the y-axis is plotted in units of proton cyclotron frequency f_{cp} , while the x-axis is plotted in units of f_{cp}/V_A where V_A is the Alfvén speed. The value of V_A differs significantly between the simulations because it is inversely proportional to the square root of the majority ion (deuteron) mass density, and hence to n_e . The




Figure 4.15: Spatio-temporal Fourier transform of the E_x field component of the simulation which makes up the left panels of Fig. 4.14, in which $n_e = 2.4 \times 10^{19} \text{m}^{-3}$. The y-axis is plotted in units of f_{cp} while the x-axis is plotted in units of f_{cp}/V_a . The horizontal black line denotes the lower hybrid frequency f_{LH} for each simulation.

horizontal black line denotes f_{LH} , below which we see excitation of the fast Alfvén wave with resonances at consecutive proton cyclotron harmonics. For frequencies less than but approaching f_{LH} , the waves gain an electrostatic component, as is evidenced by Fig. 4.15, which shows the spatio-temporal Fourier transform of the E_x field component of the simulation that makes up the left column of Fig. 4.14 in which $n_e = 2.4 \times 10^{19} \text{m}^{-3}$. Above f_{LH} there are several weaker but significant spectrally intense regions. The location of these regions in (k, f) space is at positions (k_3, f_3) such that $k_3 = k_1 + k_2$ and $f_3 = f_1 + f_2$ where (k_1, f_1) and (k_2, f_2) are the locations of strong resonances on the fast Alfvén branch below f_{LH} . We also note that the most dominant nonlinear spectral features above $f = f_{LH}$ move to increasingly high values of normalised k as density increases.

If the spectrally dense regions with co-ordinates (k_3, f_3) above f_{LH} are indeed the result of wave-wave coupling between modes below f_{LH} , this should be borne

out by bicoherence analysis of the simulated field component B_z . The corresponding bicoherence plot for each simulation is shown in the upper panels in Fig. 4.14. These plots show clearly defined sets of (k_1, k_2) pairs which have strong coupling, the most striking of which are near the $k_1 = k_2$ (and hence $f_1 = f_2$) boundary. These are modes close to each other in k space on the fast Alfvén branch. If we pick a region of strong coupling near the $k_1 = k_2$ boundary for the upper leftmost panel, say $k_1 \approx 15f_{cp}/V_A$ and $k_2 \approx 18f_{cp}/V_A$, and read off the corresponding $f_1 \approx 12f_{cp}$ and $f_2 \approx 14f_{cp}$, then we should be able to see a spectrally dense region at $k_3 \approx 33f_{cp}/V_A$ and $f_3 \approx 26f_{cp}$ in the lower leftmost plot above the $f = f_{LH}$ line. This is indeed the case, and a similar correspondence is seen across all panels of Fig. 4.14.

Bicoherence analysis of both experimental data (Fig. 4.13) and simulation outputs (Fig. 4.14) thus demonstrates strong coupling between modes near the $f_1 = f_2$ boundary below f_{LH} . This supports our conjecture that nonlinear coupling is responsible for the faint spectral “ghost” feature in Fig. 4.1, since this is also captured by our simulations. This lends further credence to our interpretation in Ref. [Chapman et al., 2017] that the downward ICE chirping is due to declining local plasma density, which is perhaps associated with the motion of an ELM filament.

4.3.2 Density dependence of downward chirping

Let us now investigate in greater depth the hypothesis that the local decline of density on submicrosecond timescales may be responsible for the downward chirping characteristics of the “ghost” ICE feature in Fig. 4.1. Due to the abundance of waves in the simulation there are many spectrally dense regions with $f > f_{LH}$ in Fig. 4.14. Accordingly, we adapt and extend the technique which was previously applied in Ref. [Chapman et al., 2017] to ICE chirping at frequencies less than f_{LH} in KSTAR. Key to this approach is analysis of the spectral properties of multiple PIC simulations, each of which is run into the nonlinear regime of the MCI at different, fixed, neighbouring values of n_e .

1. Using the experimental bicoherence plot (Fig. 4.13) along with the experimental spectrogram (Fig. 4.12), we identify spectral features “ f_1 ” and “ f_2 ” with $f < f_{LH}$, that are able to combine to produce the faint spectral features “ f_3 ” with $f > f_{LH}$.
2. The simulation with number density n_e , which in Sec. 4.1 was found to give rise to strong spectral features with frequencies f_1 and f_2 , see the left panel

of Fig. 4.16, is examined. In cases where f_1 and f_2 are present across a range of n_e values, the procedure is repeated for each simulation.

3. From this simulation, the range of values of k_1 and k_2 corresponding to f_1 and f_2 is selected. We refer to a range of values because of the finite resolution of the simulation in both wavevector and frequency space.
4. This range of k_1 and k_2 then defines the minimum and maximum values of k_3 which correspond to the f_3 observed in Fig. 4.12.
5. The spectral power as a function of frequency shown in a given vertical strip in the lower right panel of Fig. 4.16 is then obtained by integrating and averaging between the minimum and maximum possible values of k_3 .
6. As there is a one-to-one mapping between k_1 and f_1 , and between k_2 and f_2 , there is an approximate one-to-one mapping between k_3 and f_3 . Therefore the spectral power in k_3 corresponds to the power in the vicinity of f_3 .

Fig. 4.4 is reproduced here as the left set of panels in Fig. 4.16. In the lower panels of Fig. 4.16, the spectral power in the output of multiple simulations is plotted as a function of frequency and n_e , and compared with the experimental RF spectrum (upper left panel). If the faint chirping features in Fig. 4.12 are a result of wave-wave interactions between modes with $f < f_{LH}$, driven by the MCI at different densities, we expect the spectral power of the newly formed modes with $f > f_{LH}$ to exhibit a similar dependence on frequency and electron number density. To this end, the power in these modes for each simulation has been calculated, and the results are shown in the right panels of Fig. 4.16. The procedure is as follows:

The lower left and lower right panels in Fig. 4.16 have much in common. First, in each case the dominant spectral features of the simulations chirp down in frequency as electron number density decreases. Second, the density values over which this occurs declines from the pre-crash pedestal density to much smaller values, in both cases.

4.4 Steady state deuterium NBI ICE

In addition to the proton chirping ICE during KSTAR ELM crashes discussed in this chapter, “steady state” ICE consisting of spectral lines at multiple harmonics of the deuteron ion cyclotron frequency at the outer midplane edge pedestal of the KSTAR plasma, at times $\sim 10\mu s$ to $100\mu s$ before the ELM crash, are also observed on KSTAR. These two phases can be seen in Fig. 4.17, which has been adapted




Figure 4.16: Left panels: reproduction of Fig. 4.4. Right, top panel: expanded view of the upper region of Fig. 4.12 in which $f \geq f_{LH}$. Right, lower panels: chirping is apparent in frequency versus number density plots for the nonlinear stage of MCI simulations. Shading indicates the \log_{10} of the spectral power in the fluctuating part of the B_z field component of each simulation.

from Fig. 2 of [Kim et al., 2018]. Time intervals A and B broadly correspond to the steady state phase, and C and D to chirping phase.

Here we turn to the phenomenon of ICE spectra from KSTAR that are characterised by deuteron, as distinct from proton, cyclotron harmonic structure. In the two KSTAR deuterium plasmas that we consider, beams of 80keV to 100keV deuterons were injected tangentially with NBI heating power in the range 2.8MW to 3.8MW. Detailed calculations of NBI ion losses suggest that some of these NBI deuterons can remain confined on trapped orbits [Kim et al., 2016]. We therefore investigate whether this small subset of NBI deuterons can give rise to ICE via the MCI in the outer midplane edge plasma. In particular we study the fast collective relaxation of these NBI deuterons by means of a direct numerical kinetic treatment using the EPOCH PIC code. The velocity-space distribution of the energetic ions is initialised as a ring-beam distribution, and in these computations, we consider only perpendicular spatial propagation ($k_{\parallel} = 0$), so that the distribution of the parallel velocity v_{\parallel} is immaterial and can be simplified to a delta function.




Figure 4.17: Upper panel: 200MHz RF (black) and D_α (red) signals around the ELM crash during KSTAR plasma 16176. Lower panel: Temporal evolution of ICE amplitude during an ELM crash in KSTAR plasma 16176. Time $t = 0$ denotes the time at which the first derivative of the RF signal is almost discontinuous, which coincides with the peak in the RF signal, the bursting phase [Kim et al., 2018]. During windows A and B, the ICE signal shows spectral peaks at successive harmonics of the deuterium cyclotron frequency. The ICE signal during windows C and D shows more complex burst phenomena, which are discussed in Refs. [Thatipamula et al., 2016; Kim et al., 2018; Chapman et al., 2017, 2018].

4.4.1 Identifying NBI deuterons in KSTAR that could relax via the MCI

In order to carry out direct numerical simulations using the PIC self-consistent kinetic approach, it is first necessary to identify a population of NBI deuterons that could give rise to ICE at the outer midplane edge. To inform this search, we first calculate analytical linear growth rates of the MCI given by Eq. 8 of Ref. [McClements et al., 1996] across a range of NBI injection energies and pitch angles $\phi = \arcsin(v_{\perp}/|v|)$. When calculating the analytical linear growth rates using this equation, it is necessary to use an angle of wave propagation that is not strictly perpendicular to the magnetic field, and we select $\theta = 89^\circ$. This very closely resembles our simulation set-up for the $v_{\parallel} \ll v_{\perp}$ cases considered in this paper, as the Doppler shift due to the finite parallel wave-number k_{\parallel} is minimal. The results are shown as a contour plot in Fig. 4.18. The colour scale represents the \log_{10} of the growth rate of the fastest growing mode in both wavevector and frequency space. We see that NBI deuterons with higher energy and pitch angle have the stronger growth rates. This is expected, since previous studies show that the MCI is more readily excited for $v_{\perp}/v_A \sim 1$, where v_A is the local Alfvén speed; and, of course, v_{\perp} increases with both energy and pitch angle.

We now calculate orbits for deuterons with an energy of 100keV and large initial pitch angles. The choice of 100keV reflects the initial energy of KSTAR NBI deuterons. If this energy is primarily perpendicular, it corresponds to $v_{\perp}/v_A \sim 0.68$ in the KSTAR outer midplane edge plasma. This choice is also helpful from a computational physics perspective: previous PIC and PIC-hybrid simulations have shown that a value of v_{\perp}/v_A close to unity allows excitation of the MCI in a feasible amount of computational time, while maintaining high signal-to-noise ratios. The deuteron orbit calculations are carried out using the CUEBIT test particle code [Hamilton et al., 2003]. Defining poloidal flux ψ such that the poloidal magnetic field is equal to $\nabla\psi \times \nabla\varphi$, where φ is toroidal angle in (R, φ, Z) cylindrical coordinates, we set ψ according to the Solov'ev solution given in Eq. 1.21 of Sec. 1.3.2. This is applicable to a plasma in which the pressure depends linearly on ψ , and RB_{φ} is uniform, B_{φ} being the toroidal magnetic field.

$$\psi = \psi_0 \left[\frac{\gamma}{8} \{ (R^2 - R_0^2)^2 - R_b^4 \} + \frac{1-\gamma}{2} R^2 Z^2 \right]. \quad (4.1)$$

We use the following parameters to approximate typical KSTAR equilibria: $R_0 = 1.8$ m, $R_b = 1.31$ m, $\gamma = 0.7$, and $\psi_0 = 0.36$ Tm⁻². Motivated by experiment [Kim et al., 2018], we first set $B_0 = 1.8$ T.

We initialise particles near the core with a range of pitch angles, and find that some of these lead to orbits that traverse the edge region while remaining confined. The results of some of these calculations are shown in Fig. 4.19. Each panel represents a different initial pitch angle, and is accompanied by an enlarged plot of the edge region. By matching deuteron cyclotron harmonic ICE spectral peak frequencies to the known spatial dependence of magnetic field strength in KSTAR, we know that the emitting region is at $R = 2.21 \pm 0.023\text{m}$ [Kim et al., 2018]. We therefore discard orbits that do not reach this far out in radius, in addition to orbits that cross the plasma boundary. In principle not all of the fast ions that cross the plasma boundary will impact on the first wall, and some will remain ionised in the plasma. While such particles could contribute in part to the ICE drive, they are likely to be much less abundant than particles within the plasma boundary, and as such are excluded in this initial study. We find that an initial pitch angle of 80 degrees results in an orbit that crosses the plasma boundary, while an initial pitch angle of 84 degrees results in an orbit which does not traverse the location of ICE emission. Initial pitch angles between these two values result in orbits that are both within the plasma boundary, and traverse the ICE emitting region for the given values of NBI injection energy and magnetic field. We also consider $B_0 = 2.27\text{T}$, for which KSTAR deuterium ICE was also observed [Thatipamula et al., 2016], and arrive at similar conclusions: only NBI deuterons with energy close to the injection energy and within a narrow range of high pitch angles, are capable of driving the MCI. Now that we have isolated the subset of NBI deuterons that is in principle capable of driving the observed ICE at deuterium harmonics, in the following section we simulate their relaxation using the EPOCH PIC code and analyse the outputs.

4.4.2 Comparison between kinetic simulations and experiment

We have carried out three PIC simulations with parameters corresponding to two KSTAR plasmas that exhibit deuterium ICE leading up to the ELM crash. These simulations use an initially uniform electron number density $n_e = 2.5 \times 10^{19}\text{m}^{-3}$, corresponding to the density at the top of the edge pedestal before the ELM crash, as inferred from Thompson scattering measurements [Thatipamula et al., 2016]. This value of n_e was the upper bound of the multiple simulations at different fixed density that were used previously [Chapman et al., 2017, 2018] to show that the frequency chirping of KSTAR proton ICE is due to the density collapse. Our simulations use 2,000 particles per cell, with over 50,000 cells so as to adequately resolve the Debye length λ_D . This means the simulations evolve the dynamics of over 10^8 computational particles with physically correct electron-to-ion mass ratios. Each




Figure 4.18: Contour plot displaying the analytical linear MCI growth rate of the fastest growing mode as a function of pitch angle and particle energy, using a \log_{10} colour scale. Motivated by experimental observations, these calculations are restricted to frequencies below the 30th deuteron cyclotron harmonic. The linear MCI growth rate is exponentially strongest for pitch angles in the range $78^\circ < \phi < 85^\circ$. For a given pitch angle, the linear MCI is strongest at higher NBI deuteron energy; the strength of this dependence increases with pitch angle. The range of pitch angles displayed reflects the range for which the CUEBIT test particle code predicts orbits which are within the plasma boundary and traverse the ICE emitting region, see also Fig. 4.19.

simulation lasts 10 deuteron gyro-periods τ_{cD} , by which time the instability is well into its nonlinear saturated regime, which is crucial for comparison with experiment [Cook et al., 2013; Carbajal et al., 2014; Chapman et al., 2018]. The temperatures of the initially Maxwellian background thermal deuterons and electrons are set to 1keV. The beam deuterons are initialised with a pitch angle of 72.4° ; this is a value in the edge region which corresponds to a core value within the range $80^\circ < \phi < 84^\circ$ identified by the CUEBIT orbit calculations in the preceding section. We denote the bulk and beam deuteron number densities by n_D and n_{NBI} respectively, and the fast ion concentration by $\xi = n_{NBI}/n_D$.




Figure 4.19: Poloidal projection of 100keV NBI deuteron orbits calculated using the CUEBIT test particle code using a Solov’ev approximation to a typical KSTAR equilibrium. Particle orbits are initialised near the core for four different pitch angles within the range $80^\circ \leq \phi \leq 84^\circ$. Plots of full orbits are labelled a)-d) with corresponding insets to the right of each panel. Panels a), b), c), and d) show orbits with initial pitch angles of 80° , 81° , 83° , and 84° respectively.

Figure 4.20 shows the time evolution of the different particle and field contributions to the energy density in a PIC simulation with $\xi = 10^{-3}$ and $B_{0z} = 1.44\text{T}$; this value corresponds to the magnetic field in the ICE emitting region of a KSTAR plasma which has central $B_0 = 1.8\text{T}$. The energy transfer between particles and fields qualitatively resembles that of previous work [Cook et al., 2013; Carbajal et al., 2014], with the minority NBI deuterons transferring their energy to the bulk plasma and to the fields on timescales of the order of several ion gyro-periods. It is interesting to note that the energy density of the electrostatic component of the fields excited by the simulation is slightly larger than that of the electromagnetic component. This approximate equipartition of energy is common when the speed

of the minority ions is significantly less than v_A [Reman et al., 2016; Dendy et al., 2017]. Confident that the energy transfer resembles that which is characteristic of the MCI, we now run two more simulations with $\xi = 10^{-2}$. This concentration is unrealistically large compared to that expected in KSTAR, and leads to an unrealistically short saturation time, but is necessary computationally in order to obtain reasonable signal-to-noise ratios. This choice is not expected to affect our conclusions, because the simulated ICE power due to the MCI has been found to scale linearly with fast particle concentration [Carbajal et al., 2017]. Increasing ξ in the simulations thus acts to shift the simulated signal above the noise, but with no significant consequences for the underlying physics. Even with this choice of ξ , the noise level for the higher, experimentally relevant, cyclotron harmonic spectral peaks that are excited in the PIC simulations still poses a challenge; this we address by using an unusually high number of computational particles. To provide a baseline which quantifies the effect of the noise, we have run two additional simulations which correspond to a thermal background plasma without the minority energetic NBI deuterons. These “background” simulations have parameter sets which are otherwise identical to their MCI counterparts, and give rise to the green traces in Fig. 4.22 below. The residual spectral structure in the green traces reflects the concentration of noise energy at normal modes in line with the fluctuation-dissipation theorem [Birdsall and Langdon, 1985].

Figure 4.21 shows the distribution of energy in the fluctuating z -component of the magnetic field, ΔB_z , across frequency-wavenumber space, obtained from a simulation identical to that of Fig. 4.20, except for the larger value of $\xi = 10^{-2}$. This plot is a spatio-temporal Fourier transform of ΔB_z over the intervals $0 \leq x \leq 50,000\lambda_D$ and $0 \leq t \leq 5\tau_{cD}$. The vertical axis of Fig. 4.21 is normalised to the deuteron cyclotron frequency f_{cD} , while the horizontal axis is normalised to f_{cD}/V_A , where V_A is the Alfvén speed. The solid blue diagonal line denotes the initial velocity of the minority NBI deuterons, and the ratio of the perpendicular component of the NBI deuteron velocity to the Alfvén speed is ~ 0.66 . This sub-Alfvénic value highlights a computational challenge associated with resolving the high harmonics that are excited by the MCI.

The phase velocity of the fast Alfvén wave is significantly greater than the perpendicular velocity of the minority NBI deuterons v_{NBI} , which is plotted as a blue diagonal line in Fig. 4.21. For a given perpendicular wave-number k_\perp , this causes a divergence between the value of $\omega = l(2\pi f_{cD}) \approx k_\perp V_A$ and the closest value of $k_\perp v_{NBI}$. The size of this divergence increases with harmonic number l . It follows that the phase velocity of MCI-excited waves with high wavenumbers, and

hence high frequency, is further from that of the fast Alfvén wave than it is for low frequency MCI excited waves; in consequence, they have much lower growth rates. The spectral power is greatest at locations in (ω, k) space where the deuteron cyclotron harmonic waves intersect the fast Alfvén wave, and these are visible at lower harmonics in Fig. 4.21. At higher wavenumbers and frequency, this effect is less striking, but can be distinguished from the background noise.

The spatiotemporal Fourier transform shown in Fig. 4.21 has been integrated across the domain $k > 0$ to obtain the power spectrum as a function of frequency. This spectrum is shown as the blue trace in Fig. 4.22 (c). Its counterpart from a simulation with $B_z = 1.84\text{T}$ which corresponds to a KSTAR plasma with central magnetic field $B_0 = 2.27\text{T}$, is shown in Fig. 4.22 (d). These power spectra constitute the PIC simulation counterparts of the experimentally measured ICE spectra. Panel (a) of Fig. 4.22 displays the power spectrum of the experimental counterpart to the PIC simulation which gives rise to panel (c): KSTAR plasma 16176, in which the spacing between deuteron cyclotron harmonics $f_{cD} \sim 11.1\text{MHz}$. In all panels, the vertical axis is plotted on a \log_{10} scale while the horizontal axis is normalised to the deuteron cyclotron frequency. We see that the deuteron cyclotron harmonic ICE spectral peaks observed in the KSTAR plasmas are all excited by the collective relaxation of the energetic deuteron population in the PIC simulations. Across all deuteron harmonics the intensity of the MCI-excited waves (blue traces) is one or two orders of magnitude higher than that of the thermal plasma noise (green traces). It can be seen that the spectral peaks in the simulations are less well resolved in frequency than in the experimentally measured signal. This is because the instability in the simulation reaches saturation so rapidly, owing to the unrealistically large value of ξ which is necessary to achieve a sufficient signal-to-noise ratio. It is encouraging to see the broadening and diminishing amplitude, of spectral peaks at $l \gtrsim 21$ in both experiment and simulation.

Panels (b) and (d) of Fig. 4.22 show similar traces corresponding to KSTAR plasma 11474 with central magnetic field $B_0 = 2.27\text{T}$ and corresponding $f_{cD} \sim 14.0\text{MHz}$ at the outer midplane edge. The agreement between PIC simulation outputs and the measured experimental ICE spectrum is good for harmonics $l \gtrsim 11$. However several spectral peaks with $l < 11$ are present in the simulation but are not observed in the experiment. We consider that these peaks may well have been excited, but were not detectable because the antenna used to measure the RF data whose power spectrum is plotted in panel (a) of Fig. 4.22 was a spiral antenna, optimised for circularly polarised waves. A different, Bowtie antenna, optimised for linearly polarised waves, was used for KSTAR plasma 11474, corresponding to panel




Figure 4.20: Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time, from a PIC simulation with $\xi = 10^{-3}$. The traces, ordered from top to bottom at their peak (and in colour online) are: Top (red) the change in kinetic energy density of the thermal bulk plasma deuterons; second (black) change in energy density of the electrons; third (blue) the energy density of the electrostatic field E_x ; fourth (green) the energy density of the magnetic field perturbation ΔB_z ; fifth (cyan) the change in kinetic energy density of the minority energetic NBI deuterons. Time is normalised to the deuteron gyro period. The primary energy flow from the NBI deuterons is to the thermal deuterons, whose kinetic oscillation helps support the field oscillations excited by the MCI. These field oscillations include, with comparable magnitude, an electromagnetic component $(\Delta B_z)^2$ and an electrostatic component E_x^2 . The electrostatic component involves electron kinetics which are fully captured in our PIC model. The MCI saturates within five deuteron gyroperiods. The initial energy densities of the E_x field component and the fluctuating ΔB_z field component are both zero; while the initial energy densities of the electrons, thermal deuterons, and NBI deuterons are $\sim 6 \times 10^3 \text{Jm}^{-3}$, $\sim 6 \times 10^3 \text{Jm}^{-3}$, and $\sim 360 \text{Jm}^{-3}$ respectively.

(b) of Fig. 4.22. The S_{11} return loss [Bird, 2009] of these two antennas is plotted in Fig. 4.23. For frequencies less than 150MHz ($l < 11$), the return loss of the Bowtie antenna is close to 0dB. This implies a high degree of reflectivity for an incoming signal, which would therefore be undetectable. Combined with the linear or circular optimisations of the two antenna polarisations, this offers a likely explanation of




Figure 4.21: Distribution of energy in the fluctuating z -component of the magnetic field ΔB_z across frequency-wavenumber space from a PIC simulation with $\xi = 10^{-2}$, $T_e = T_D = 1\text{keV}$, $B_{0z} = 1.44\text{T}$, $n_e = 2.5 \times 10^{19}\text{m}^{-3}$, and a 100keV minority NBI deuteron population. This plot is a spatio-temporal Fourier transform of the B_z field over the intervals spanning $0 \leq x \leq 50,000\lambda_D$ and $0 \leq t \leq 5\tau_{cD}$. Shading indicates the \log_{10} of the spectral density of the oscillatory part ΔB_z of the B_z field component in frequency-wavenumber space. The sweep of the fast Alfvén wave from bottom left to top right is intersected by cyclotron harmonic waves at successive deuteron harmonics. The phase velocity of the fast Alfvén wave $\simeq v_A$, and this exceeds the speed v_{NBI} of the NBI deuterons which is plotted as a blue diagonal line. Wave excitation is strongest in the wedge between v_A and v_{NBI} , and in particular where cyclotron harmonic waves intersect the fast Alfvén wave. Simulated ICE frequency spectra, such as the lower panels of Fig. 4.22, are obtained by integrating plots such as Fig. 4.21 over wavenumber.

the mismatch between experimental and simulation spectral peaks at low harmonic numbers seen in panels (b) and (d) of Fig. 4.22.




Figure 4.22: Panels (a) and (b): measured spectral intensity of the ICE signal from KSTAR deuterium plasmas with 100keV deuteron NBI heating: (a) [Kim et al., 2018] plasma 16176 with $B_0 = 1.84\text{T}$, (b) [Thatipamula et al., 2016] plasma 11474 with $B_0 = 2.27\text{T}$. In both panels the FFT was performed using data from $t = -150\mu\text{s}$ to $t = -100\mu\text{s}$ relative to t_0 (see, for example, Fig. 4.17). Panels (c) and (d): Blue traces are outputs of PIC simulations of the spectral intensity of the fluctuating B_z field energy density, resulting from relaxation of a minority 100keV deuteron ring-beam population in thermal deuterium plasma. The simulation parameters of (c) and (d) map to (a) and (b) respectively, corresponding to the local ICE-emitting plasma in the two KSTAR experiments: $n_e = 2.5 \times 10^{19}\text{m}^{-3}$, $T_e = T_D = 1\text{keV}$. Green traces provide a noise baseline for the blue traces. They are obtained from the thermal plasma without a ring-beam, so that any spectral structure arises from the fluctuation-dissipation theorem and identifies normal modes.

4.5 Conclusions

We have shown that harmonic ICE with spacing equal to f_{cp} in KSTAR deuterium plasmas can be driven by a small subset of the fusion-born proton population originating in the core of the plasma. The drift orbits of these protons have large radial excursions to the outer midplane edge. We have compared the nonlinearly saturated field spectra obtained from multiple MCI simulations at different plasma densities




Figure 4.23: The return loss (S_{11}) as a function of frequency of the Spiral (blue) and Bowtie (red) antennas used for the detection of signals in panels (a) and (b) respectively. The y-axis denotes the return loss and is on a log scale. For frequencies less than 150MHz, the return loss of the Bowtie antenna is close to 0dB, offering a likely explanation of the mismatch between experimental and simulation spectral peaks at low harmonic number ($l \leq 11$) seen in panels (b) and (d) of Fig. 4.22.

with experimentally observed time evolving ICE spectra. These results suggest that downward chirping of ICE occurs when the emitting fusion-born protons ions are embedded in a local plasma with rapidly falling density. These results also show that the much rarer upward chirping of ICE occurs when the local density is rapidly rising. By combining different simulation spectra with the chirping ICE observed during KSTAR ELM crashes, we have obtained sub-microsecond time resolution measurements of the evolving electron number density in the emitting region. We conjecture that the time-evolution of the electron number density on such rapid timescales in this region results from the motion of an ELM filament; for a detailed account of the experimental set-up and the detection of chirping ICE in conjunction with ELM filaments, see Refs. [Thatipamula et al., 2016; Kim et al., 2018].

The “ghost” ICE feature in Fig. 4.12 is a real plasma physics phenomenon. Its existence is due to a combination of energetic particle physics with linear and nonlinear wave physics, which is so far observed only in KSTAR tokamak plasmas. Here we have shown that the separate, fainter (“ghost”) chirping ICE feature observed in the frequency range 500 MHz ($20f_{cp}$) to 900 MHz ($36f_{cp}$) is driven by nonlinear wave coupling between different neighbouring cyclotron harmonic peaks in the main ICE feature below 500 MHz. This is evident from bispectral analysis of: first, the measured KSTAR fields, where we benefit from exceptionally high (up to 20 GS s^{-1}) sampling rates; and second, field amplitudes output from first

principles particle-in-cell code simulations of the KSTAR fusion-born proton relaxation scenario. This reinforces the MCI interpretation of chirping proton ICE in KSTAR. It also provides a novel demonstration of nonlinear wave coupling on very fast timescales in a tokamak plasma.

By combining the linear analytical theory of the MCI, energetic particle orbit studies, and first principles PIC simulations, we have been able to provide an explanation for the origin of steady-state ICE at multiple deuterium cyclotron harmonics, which is observed in KSTAR deuterium plasmas heated by deuterium NBI. We first identified a small subset of the NBI deuterium population that could be responsible for the emission. This was done by performing test particle calculations for deuterons with a range of NBI-relevant energies and pitch angles. We identified those that are confined, pass through the outer midplane edge region where the ICE originates, and have the largest analytical MCI growth rates. We then carried out two high resolution PIC simulations with parameters corresponding to the ICE emitting region at the outer mid-plane edge of two KSTAR plasmas. The initial velocity distributions of the kinetic thermal deuterons and electrons in these PIC simulations is Maxwellian. In addition, there is an initial minority energetic deuterium population, whose velocity distribution reflects NBI parameters and our orbit studies. The collective relaxation of the NBI deuterium population in these two PIC simulations generates electric and magnetic field oscillations whose power spectra substantially resemble the measured ICE spectra. Some low harmonic peaks in one simulation frequency spectrum were not detected in its experimental counterpart. A probable explanation for this is that the S_{11} return loss of the Bowtie antenna used to measure this RF signal was close to 0dB, implying very high reflectivity in this low frequency range.

We have now explained the main features of two distinct types of ICE from KSTAR plasmas: steady-state ICE due to NBI deuterons; and highly dynamic chirping ICE due to fusion-born protons. We note that upward and downward chirping ICE with spacing f_{cD} has been observed in Refs. [Thatipamula et al., 2016; Kim et al., 2018], and this is a phenomenon under investigation.

This work opens unexpected lines of research linking fusion-born ion physics with ELM crashes in medium size tokamaks such as KSTAR, which generated the ICE observations studied here. It also provides a new, very high time resolution, method of quantifying the experimental phenomenology. The discovery that one can

investigate collective fusion-born ion physics in medium size tokamaks, the connection to ELM phenomenology, and the high time resolution, are all new. Because of the relatively wide access that exists to medium-size tokamaks worldwide, combined with the centrality of fusion reactivity and ELMs to tokamak plasma physics, wider exploitation of these results should be possible. The successful interpretation of this unexpected phenomenon spontaneously driven by fusion-born ions, helps establish interpretive capability for ICE from future deuterium-tritium plasmas in JET and ITER. Modelling of the plasma physics underlying ICE signals yields information on two key features of the ICE-emitting energetic ion population. First, the values of key parameters, notably the ratio of the characteristic perpendicular velocity of the energetic ions to the local Alfvén speed. This needs to be of order unity. Second, the structure of the distribution of the energetic ions in velocity space, which needs to be strongly non-Maxwellian in order to excite the MCI which underlies ICE. A drifting ring distribution, i.e., the product of two delta-functions in parallel and perpendicular velocity, has been found to be the best few-parameter way of capturing this structure for ICE applications. This approximation has proven fruitful across more than two decades, spanning ICE measurements from deuterium-tritium plasmas in JET [Cottrell and Dendy, 1988] and TFTR [Cauffman et al., 1995] during the mid-1990s to the most recent measurements reported from ASDEX-Upgrade in 2014 [D’Inca, 2014] and JT-60U in 2017 [Sumida et al., 2017, 2018]. The new results presented here confirm the fidelity of the output of first principles PIC simulations in relation to measured ICE signals, alongside the validity of the model for ICE that is implemented in the PIC code. The agreement between the bispectral analysis of the simulation outputs and the observations of an unexpected, strongly nonlinear, transient ICE feature provides fresh validation of the ICE model, in a challenging regime. The ICE physics addressed here would have gone unnoticed had it not been for KSTAR’s sophisticated RF system and high speed digitizer [Leem et al., 2012]. Their uniquely high time resolution has yielded new insights into the dynamics of ELMs, energetic ions, and wave phenomena in tokamak plasmas. The present results further strengthen the case for the use of ICE as a diagnostic [McClements et al., 2015; Dendy and McClements, 2015; Gorelenkov, 2016] on ITER [Hawryluk et al., 2009; Kessel et al., 2009], from both an operational and a fundamental physics perspective.

Chapter 5

Simulations of ICE in JET and ASDEX Upgrade

In this chapter we discuss simulations of the MCI relevant to two tokamaks. Firstly, we consider recent JET results [Jacquet et al., 2011; McClements et al., 2018] in which ICRF ICE was detected at the edge region of the plasma corresponding to the first harmonic of ionized ^3He . We then discuss the possibility of the MCI being the driving mechanism behind the core ICE reported in Ref. [Ochoukov et al., 2018], in which emission was observed at the fundamental proton/second deuterium harmonic.

5.1 JET

In Refs. [Jacquet et al., 2011; McClements et al., 2018], measurements of ICE in JET plasmas heated by waves in the ion cyclotron range of frequencies were reported. Hydrogen was the majority ion species in all of the plasma pulses and the measurements were obtained using ICRF antennas. In most cases the energetic ions can be clearly identified as ICRF wave-accelerated minority ^3He ions, and orbit calculations using the CUEBIT code [Hamilton et al., 2003] support the hypothesis that the observed emission at the fundamental ^3He cyclotron harmonic was produced by energetic ^3He ions undergoing drift orbit excursions to the outer mid-plane plasma edge. In this section we perform multiple simulations using plasma parameters corresponding to edge plasma conditions in these JET pulses, and energetic particle parameters inferred from the cyclotron resonance location. In all of the following simulations, the temperature of the background protons and electrons $T_p = T_e = 100\text{eV}$, and the electron number density $n_e = 1 \times 10^{19}\text{m}^{-3}$, consistent

with Thompson scattering measurements [McClements et al., 2018]. The magnetic field is set to 1.995T consistent with experiment, and the minority ions are represented by a ring beam distribution with the ratio of minority ${}^3\text{He}$ ions to majority protons $\xi = n_{3\text{He}}/n_p = 10^{-3}$.

In Ref. [McClements et al., 2018], the parallel and perpendicular velocities of the 2MeV minority ${}^3\text{He}$ ions were calculated to be $v_{\parallel} = 5 \times 10^6 \text{ms}^{-1}$ and $v_{\perp} = 1 \times 10^7 \text{ms}^{-1}$ respectively. We first run three simulations to examine the effect of a parallel velocity drift in the minority ${}^3\text{He}$ distribution. One simulation uses a wave propagation angle of $\theta = 90^\circ$, and an initial parallel velocity $v_{\parallel 0} = 0$, as the parallel velocity is not expected to play a role for strictly perpendicular propagation. The other two simulations use $\theta = 89^\circ$, one with $v_{\parallel 0} = 0$, and another with $v_{\parallel 0} = 5 \times 10^6 \text{ms}^{-1}$. In all three of these simulations, and subsequent simulations in this section, the initial perpendicular velocity of the minority ${}^3\text{He}$ ions is set to $v_{\perp 0} = 1 \times 10^7 \text{ms}^{-1}$. These three simulations use 70,000 grid cells each resolving the Debye length and with the number of particles per cell set to 100. The simulations were run for 9 He^3 gyro periods $\tau_{g\text{He}^3}$, well into the nonlinear regime. Only the first six minority ${}^3\text{He}$ gyroperiods of the simulation are shown in the following plots of energy density as a function of time. Beyond $6\tau_{g\text{He}^3}$, there is little change in the energy transfer, but the extended data set allows for a larger Fourier transform window and hence better resolved power spectra and dispersion relations.

Figure 5.1 plots the change in energy density of two simulations with $\theta = 90^\circ$ on the left, and $\theta = 89^\circ$ on the right. In both panels, we see the characteristic MCI energy exchange between the minority ${}^3\text{He}$ ions, the background plasma, and the electromagnetic fields. The linear stage of the MCI instability saturates quickly at around $3\tau_{g\text{He}^3}$, with the $\theta = 90^\circ$ simulation taking slightly longer to reach linear saturation. In the linear stage, the amount of energy transfer for the minority ions to the majority ions is larger in the $\theta = 90^\circ$ than in the $\theta = 89^\circ$. In the nonlinear stage of the simulation, the amount of re-energisation [Carbajal et al., 2014] of the minority ions is noticeably less for $\theta = 89^\circ$ than in the $\theta = 90^\circ$ simulation, which may influence the character of the nonlinearly driven spectral features present in the simulations.




Figure 5.1: Left: time evolution of the change in energy density of particles and electric and magnetic fields as a function of time for a PIC simulation of the MCI with JET like plasma parameters and a wave propagation angle of $\theta = 90^\circ$. Right: as left but using $\theta = 89^\circ$. In both cases the parallel component of the ^3He ring beam velocity $v_{\parallel 0}$ is set to 0 and time is normalised to the ^3He gyro-period.

The dispersion relation of the z-component of the magnetic field in these simulations is plotted in Fig. 5.2 and the spectral intensity as a function of frequency is plotted in Fig. 5.3. The temporal Fourier transform is calculated between $0\tau_{g\text{He}^3} \leq t \leq 9\tau_{g\text{He}^3}$, well into the nonlinear regime. The spatial Fourier transform is calculated across the entire simulation domain. The dispersion relation of the two simulations look very similar, a noticeable feature being the strong low frequency nonlinear feature below the third cyclotron harmonic spanning full range of displayed wavevectors. In the right panel of Fig. 5.2, this low frequency feature has a strong resonance with the third harmonic. Fig. 5.3 shows that the fundamental harmonic is excited in both cases, consistent with experimental observations. Both panels also show excitation of harmonics above the fundamental, but the low l harmonics are sharper in the left panel, while the high l harmonics are sharper in the right panel. The right panel of Fig. 5.3 shows a more intense third harmonic than that of the left panel, which is due to the aforementioned resonance shown in the right panel of Fig. 5.2.

Figure 5.4 compares the time evolution of the energy density between the two $\theta = 89^\circ$ simulations. The left panel is from a simulation with $v_{\parallel 0} = 0$, while the right panel is from a simulation with $v_{\parallel 0} = 5 \times 10^6 \text{ms}^{-1}$. The right panel with a finite minority ion parallel velocity reaches linear saturation slightly later than the $v_{\parallel 0} = 0$ simulation. The only other noticeable difference is the degree of oscillation between the fields and particles, with the right panel being less oscillatory than the




Figure 5.2: Distribution of energy in the B_z field component across frequency-wavenumber space. Shading indicates the \log_{10} of the spectral density. The temporal Fourier transform is calculated between $0\tau_{g\text{He}^3} \leq t \leq 9\tau_{g\text{He}^3}$, while the spatial Fourier transform is calculated across the entire simulation domain. The vertical axis is normalised to the ${}^3\text{He}$ cyclotron frequency while the horizontal axis is normalised to the ${}^3\text{He}$ cyclotron frequency divided by the Alfvén speed. Left: $\theta = 90^\circ$. Right: $\theta = 89^\circ$.




Figure 5.3: Spectral intensity of the B_z field component of the MCI simulation using a temporal Fourier transform window spanning data from $0\tau_{gp} \leq t \leq 9\tau_{gp}$. The vertical axis is on a \log_{10} scale while the horizontal axis is normalised to the ${}^3\text{He}$ cyclotron frequency. Left: $\theta = 90^\circ$. Right: $\theta = 89^\circ$.

left, but the overall levels of energy transfer are almost identical at all points in the simulation. The dispersion relations corresponding to these two simulations are shown in Fig. 5.5, and the power spectrum using data Fourier transformed between $0\tau_{gp} \leq t \leq 9\tau_{gp}$ is shown in Fig. 5.6. We can see from the right panel of Fig. 5.6 that the spectral structure is slightly clearer than that of its $v_{\parallel 0} = 0$ counterpart,

and overall the intensity of the instability is reduced, with the exception of the $l = 3$ and $l = 8$ harmonics.




Figure 5.4: Left: time evolution of the change in energy density of particles and electric and magnetic fields as a function of time for a PIC simulation of the MCI with JET like plasma parameters and $\theta = 89^\circ$. The parallel component of the ring beam velocity $v_{\parallel 0} = 0$ and time is normalised to the ${}^3\text{He}$ gyro-period. Right: as left except $v_{\parallel 0} = 5 \times 10^6 \text{ms}^{-1}$.




Figure 5.5: Spectral intensity of the B_z field component of the MCI simulation with $\theta = 89^\circ$ and using a temporal Fourier transform window spanning data from $0\tau_{gp} \leq t \leq 9\tau_{gp}$. Left: the parallel component of the ${}^3\text{He}$ ring beam velocity $v_{\parallel 0} = 0$. Right: $v_{\parallel 0} = 5 \times 10^6 \text{ms}^{-1}$.




Figure 5.6: Spectral intensity of the B_z field component of the MCI simulation with $\theta = 89^\circ$ and using a temporal Fourier transform window spanning data from $0\tau_{gp} \leq t \leq 9\tau_{gp}$. The vertical axis is on a \log_{10} scale while the horizontal axis is normalised to the ${}^3\text{He}$ cyclotron frequency. Left: the parallel component of the ${}^3\text{He}$ ring beam velocity $v_{\parallel 0} = 0$. Right: $v_{\parallel 0} = 5 \times 10^6 \text{ms}^{-1}$.

We now run four additional simulations with wave propagation angles of $90^\circ, 92^\circ, 94^\circ$ and 96° with respect to the simulation domain, each with $v_{\parallel 0} = 0$. Due to the symmetrical nature of our ring-beam distribution for the minority ions (and the Maxwellian distributions for the thermal species), results from simulations using $\theta = 92^\circ, 94^\circ$ and 96° are the same as those from simulations using $\theta = 88^\circ, 86^\circ$ and 84° respectively. The simulation set-up is the same as the other simulations in this section, apart from the number of particles per cell which has been increased to 900, thus reducing the numerical noise levels in the simulations by a factor of $\sqrt{900/100} = 3$. The energy dynamics of these simulations are shown in Fig. 5.7, in which panel a) is nearly identical to the simulation displayed in the left panel of Fig. 5.1, differing only in the number of particles per cell. Reassuringly, panel a) of Fig. 5.7 is qualitatively the same as the left panel of Fig. 5.1, with the linear stage of the former lasting slightly longer than that of the latter, most likely due to the reduced particle noise which acts to seed the initial linear growth. As the angle of wave propagation θ increases to 92° , the amount of energy transfer between the minority ions and the background ions and fields decreases by $\sim 50\%$. Interestingly, the majority of the energy transfer is to the x-component of the electric field, as well as a large portion to the electrons, suggesting the instability is predominately electrostatic in this regime. At $\theta = 94^\circ$, the MCI is barely excited, if at all, with almost all the energy from the minority ions going directly into the thermal electrons. At $\theta = 96^\circ$, the instability is most certainly not excited, the thermal electrons are

gradually extracting all the energy from the minority ions.

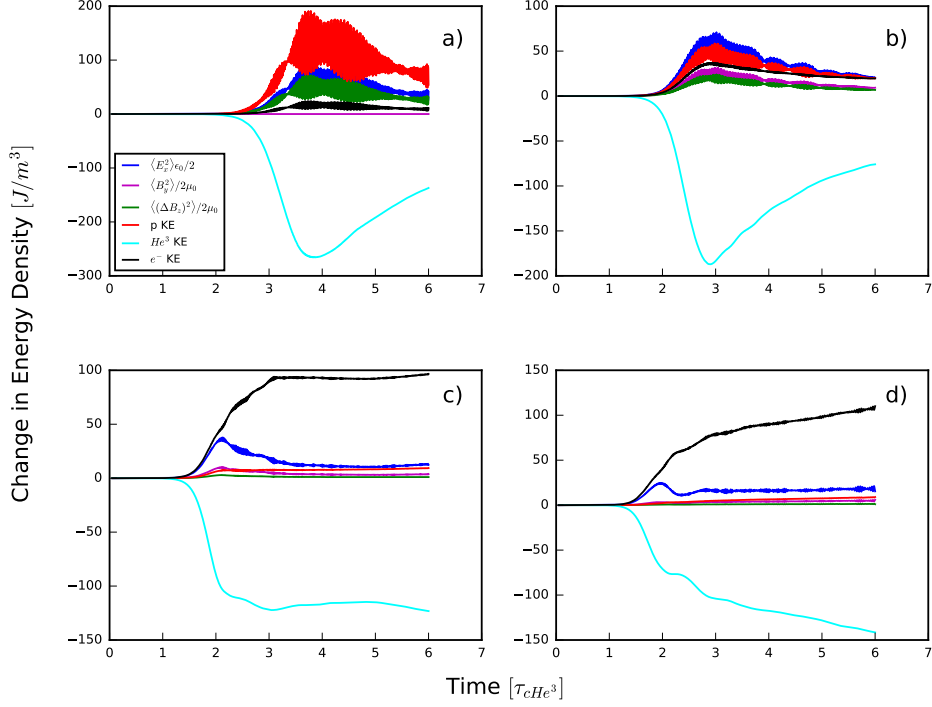


Figure 5.7: Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time for four PIC simulations of the MCI with JET like plasma parameters. Panels a), b), c) and d) correspond to simulations with wave propagation angles equal to 90° , 92° , 94° and 96° respectively. In each panel the parallel component of the ring beam velocity $v_{\parallel 0}$ is set to 0 and time is normalised to the ^3He gyro-period.

5.2 ASDEX Upgrade

As discussed in Sec. 1.4, a variety of ion cyclotron phenomena have been observed in the ASDEX-Upgrade (AUG) tokamak, including FP ICE from both the core and edge regions, NBI ICE, and ICRF ICE [D’Inca, 2014; Ochoukov et al., 2018, 2019]. In this section, we discuss the possibility of the MCI being the driving mechanism behind the core ICE reported in Ref. [Ochoukov et al., 2018], in which emission was observed at the fundamental proton/second deuterium harmonic. The only protons in these pulses were those due to fusion reactions, while the energetic deuterons were due to neutral beam injection (NBI). The Alfvén velocity in the core AUG plasma was $5-6 \times 10^6 \text{m}^{-3}$, which, while below the proton birth velocity ($2.4 \times 10^7 \text{ms}^{-1}$), is a factor of two above the NBI deuteron velocity ($2.4 \times 10^6 \text{ms}^{-1}$). Previous studies have shown that ICE is more readily excited when the energetic ions are super-Alfvénic. For this reason, it was concluded that the most likely driver of the ICE emission is fusion-born protons. As discussed in previous sections, a key parameter for MCI excitation is the ratio of the perpendicular component of the energetic ion velocity $v_{\perp 0}$ to the local Alfvén velocity v_A . Defining the perpendicular component of the energetic proton velocity as $v_{\perp 0}$, if we assume equipartition of energy, we have $v_{\perp 0} = 1.96 \times 10^7 \text{ms}^{-1}$. Using the upper bound of the estimate for v_A in the core, we have $v_{\perp 0}/v_A \sim 3.26$. Using this value of $v_{\perp 0}$, a PIC simulation of the MCI in which the background plasma was that of a typical AUG plasma was carried out. A thin spherical shell distribution was used for the energetic proton ions, as this closely resembles the actual particle distribution in the experiment, that of freshly born fusion ions in the core. After running the simulation for more than 120 proton gyro-periods, the MCI was not excited and there was no energy transfer between any of the particle species and fields. As discussed in chapters 2 and 3, PIC simulations essentially approximate particle phase space by a series of delta functions in the form of computational macro particles or “markers”. It seems reasonable to assume that as ICE is observed in AUG experiments, and the distribution function of the fusion-born protons in the simulations closely resembles that of the experiment, ICE should be excited in the simulation, provided the phase space of the energetic protons is adequately resolved. It is possible that by modelling the entire shell distribution functions of energetic protons, we are resolving huge areas of phase space with $v_{\perp 0}/v_A \gg 1$ that do not contribute to the excitation of the MCI. To test this, one could sample the areas of phase space with $v_{\perp 0}/v_A \sim 1$ to a given tolerance using more computational macro particles, thereby more accurately representing the energetic proton distribution function in this region of phase space.

As a first approximation to this, we suppose that a thin spherical shell distribution can be approximated by multiple thin ring distributions which span the perpendicular plane in velocity space, and are stacked on top of one another such that their position in the stack denotes the parallel component of their velocity. If we stack an infinite amount of these rings on top of one another, making their radius infinitely small at the extrema of the stack, we recover our shell distribution. In the discretised version of this picture, only a select few of these rings satisfy $v_{\perp 0}/v_A \sim 1$. In this spirit, we have run another simulation of the MCI using the same AUG plasma parameters, however, this simulation uses a ring beam distribution to represent the minority energetic protons, with $v_{\perp 0} \simeq 7.58 \times 10^6 \text{ms}^{-1}$, giving $v_{\perp 0}/v_A \sim 1.16$. The simulation uses a background magnetic field $B_z = 2.62\text{T}$ oriented at 89° to the simulation domain. The parallel component of the energetic proton velocity $v_{\parallel 0}$ is neglected as it does not contribute to the largely perpendicular wave propagation. To aid with frequency resolution, a 10% perpendicular velocity spread is introduced. This has the effect of decreasing the linear growth rate of linearly unstable modes, resulting in a slightly longer simulation time and slightly better frequency resolution. The thermal plasma comprises deuterons and electrons, with temperature $T_D = T_e = 5\text{keV}$. The electron number density $n_e = 3.8 \times 10^{19} \text{m}^{-3}$, and the ratio of energetic to background ions $\xi = n_p/n_D = 10^{-3}$. In the discharge considered here, see Fig. 5 of Ref. [Ochoukov et al., 2018], only the fundamental proton cyclotron harmonic at around $f_{cp} = 40\text{MHz}$ was observed as the frequency response rate of the B-dot probe used in the ICE measurements is limited to $10 - 50\text{MHz}$ [Ochoukov et al., 2015]. Resolving the fundamental harmonic in PIC simulations can be difficult, as the fundamental, along with other harmonics with $l \lesssim 5$, are often driven exclusively by non-linear interactions among MCI excited modes with $l > 5$ [Carbajal et al., 2014; Carbajal, 2015]. In simulations of the MCI, this nonlinear drive, while essential to understand the full extent of ICE in simulations and experiment, is usually weaker than the linearly unstable modes in the simulation. This, combined with the particle noise present in the low frequency region of PIC simulations, means that the PIC simulation here must be incredibly well resolved. In this simulation, a total of 38,000 grid cells and 4,000 particles per cell are used while resolving the Debye length and using higher order particle weighting functions. The simulation is run for 20 proton gyro-periods τ_{gp} , reaching linear saturation by $t \approx 10\tau_{gp}$. To quantify the effect of particle noise, a companion simulation is performed without the minority energetic protons.

As a first check, the energy transfer in the MCI simulation is plotted in Fig. 5.8. This looks very much like the “standard” ICE simulations shown in chapter

3, the minority protons transfer their energy predominantly to the z-component of the background magnetic field, as well as the thermal deuterons. The simulation takes almost 10 gyro-periods to reach linear saturation, after which the nonlinear, re-energisation stages takes over the dynamics [Carbajal et al., 2014; Carbajal, 2015].




Figure 5.8: Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time. The traces, ordered from top to bottom at their peak are: Top (red) the change in kinetic energy density of the thermal bulk plasma deuterons; second (green) the energy density of the magnetic field perturbation ΔB_z ; third (blue) the energy density of the electrostatic field E_x ; fourth (magenta) the energy density of the magnetic field B_y ; fifth (cyan) the change in kinetic energy density of the minority energetic protons. Time is normalised to the proton gyro-period. The MCI saturates within ten proton gyro-periods.

The power spectra of the B_z field component as a function of frequency ω is shown in Fig. 5.9 and is plotted on a \log_{10} scale. This plot is obtained by a temporal Fourier transform of the first 10 gyro-periods of the simulation, so just into the nonlinear phase, and integrating across the entire spatial domain. Here, the blue trace denotes the power in the simulation of the MCI, while the green trace denotes the power in the background thermal plasma. Energetic proton cyclotron

harmonics are easily excited in the range $7 \leq l \leq 15$, with the strongest harmonics at $l = 11$ and $l = 12$ being more than five orders of magnitude more powerful than the background plasma signal. Modes with $l \leq 6$ appear to be linearly stable. The Fourier power as a function of time t and the wavevector k is plotted in Fig. 5.10. For (approximately) perpendicular wave propagation in the ion cyclotron range of frequencies, we benefit from an approximately one-to-one mapping between wavevector and frequency, meaning Fig. 5.10 yields information about the spectral content in the proton cyclotron harmonics as a function of time. Modes with $k \geq 8$ are clearly linearly unstable, and are the most dominant modes in the simulation. These correspond to the peaks with $\omega \geq 8$ in Fig. 5.9. There are three more striking features of this plot. These are nonlinearly driven modes at $k \simeq 1.2, 2.5,$ and 7.1 , which correspond to $\omega \simeq 1, 2,$ and 7 . These modes are purely non-linearly driven, with the mode at $(k, \omega) \approx (7.1, 7)$ being clearly visible in Fig. 5.9, meaning this mode was “switched on” very soon after the end of the linear phase.




Figure 5.9: Blue trace: spectral intensity of the B_z field component of the MCI simulation using a temporal Fourier transform window spanning data from $0\tau_{gp} \leq t \leq 10\tau_{gp}$. The green trace shows the power in a background thermal plasma without a minority energetic proton ring beam. Power is obtained by integrating across the entire spatial domain and is plotted on a \log_{10} scale.




Figure 5.10: Spectral intensity of the B_z field component of the MCI simulation as a function of time and wavevector plotted on a \log_{10} scale. Linearly unstable modes with $k > 8$ can be seen evolving in time. Three nonlinearly driven modes at $k \simeq 1.2$, 2.5, and 7.1, which correspond to $\omega \simeq 1$, 2, and 7 can be seen to evolve from the nonlinear stage of the simulation at $t \approx 9\tau_{gp}$.

A spatio-temporal Fourier transform calculated between $0\tau_{gp} \leq t \leq 18\tau_{gp}$ and across the entire simulation domain is shown in Fig. 5.11. Both forward and backward propagating waves are plotted. The left panel shows a wide range of frequencies and wavenumbers, and a wealth of nonlinear interactions manifesting themselves as orange spectrally dense regions close, but not quite along, the fast Alfvén branch and spread across a wide range of Fourier space. The right panel of Fig. 5.11 magnifies the low frequency modes in the simulation. We can see the two nonlinear modes at $k = 1$ and $k = 2.5$ lie just to the right (left) of the forward (backward) propagating fast Alfvén wave.

We now plot the power spectra as a function of frequency in the low frequency, linearly stable modes for Fourier transform windows of different duration. This is shown in Fig. 5.12, in which panels a), b), c), and d) correspond to temporal Fourier transform windows of duration $10\tau_{gp}$, $12.5\tau_{gp}$, $15\tau_{gp}$, and $17.5\tau_{gp}$ respectively. It is




Figure 5.11: Distribution of energy in the z-component of the magnetic field across frequency-wavenumber space. Shading indicates the \log_{10} of the spectral density. The temporal Fourier transform is calculated between $0\tau_{gp} \leq t \leq 18\tau_{gp}$, while the spatial Fourier transform is calculated across the entire simulation domain. The vertical axis is normalised to the proton cyclotron frequency while the horizontal axis is normalised to the proton cyclotron frequency divided by the Alfvén speed. Left: forward and backward propagating waves across a wide range of frequencies and wavenumbers. Right: an expanded view of the low frequency region of the dispersion plot. We see regions of high spectral density at proton cyclotron harmonics lying just to the right (left) of the forward (backward) propagating fast Alfvén wave.

clear that proton cyclotron harmonics $l = 1$ and $l = 2$ increase in power as the simulation progresses deeper into the nonlinear stage, with $l = 1$ showing the most pronounced effect. The Fourier power in proton cyclotron harmonics with $l > 2$ does not change substantially as the simulation progresses. It should be noted that the power in the background plasma is increasing slightly as the simulation progresses, but this increase is approximately constant and consistent across the whole range of frequencies (including proton cyclotron harmonics with $l > 6$ not shown here) for the entire duration of the simulation. In the linear stage of the simulation, any increase in Fourier power in modes with $l \leq 6$ exactly follows that of the background thermal plasma simulation, and is therefore due only to thermal noise. Due to the excellent resolution of the simulation, we can also see structure in the deuteron cyclotron harmonics of the background plasma, half of which are degenerate with the proton cyclotron harmonics. The spectral structure seen here reflects the concentration of noise power at normal modes through the fluctuation dissipation theorem. To further convince ourselves that the observed spectral power at $l = 1$ is a result of nonlinear interactions between MCI excited modes of higher frequency, we plot the evolution in the difference between MCI and thermal plasma Fourier power as a function of time. This is shown in Fig. 5.13 starting from $t = 7\tau_{gp}$

up to $t = 17.5\tau_{gp}$. We observe a gradual increase from $t = 7\tau_{gp}$ to $t = 9\tau_{gp}$, followed by a small decrease occurring at $t = 9.5\tau_{gp}$, around the transition from the linear to nonlinear regime. After this, the difference in power remains almost constant until $t = 12.5\tau_{gp}$, after which it is followed by an approximately linear increase until $t = 17.5\tau_{gp}$, by which point the power in the $l = 1$ proton cyclotron harmonic is almost an order of magnitude greater than that of the background.

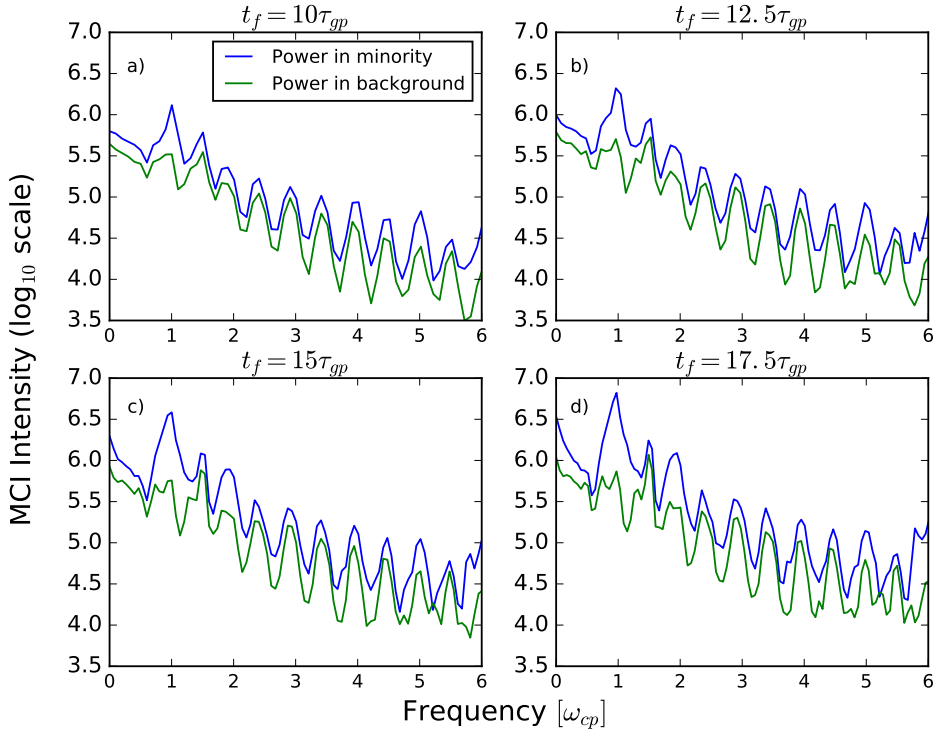


Figure 5.12: Spectral intensity of the B_z field component of the MCI simulation using temporal Fourier transform windows of increasing duration. Panels a), b), c), and d) correspond to temporal Fourier transform windows of duration $10\tau_{gp}$, $12.5\tau_{gp}$, $15\tau_{gp}$, and $17.5\tau_{gp}$ respectively. All panels show data corresponding to the nonlinear stage of the simulation. The green trace shows the power in a background thermal plasma without a minority energetic proton ring beam, and hence without the excitation of the MCI. Power is obtained by integrating across the entire spatial domain and is plotted on a \log_{10} scale.

To satisfy the frequency and wavenumber matching conditions discussed in Sec. 2.2, the linearly excited waves responsible for producing the nonlinear feature at $\omega = 1\omega_{cp}$ must be between waves travelling in opposite directions. Looking closely at Fig. 5.11, we see that there is not an exact symmetry between forward




Figure 5.13: The difference in Fourier power of the B_z field component between the $\omega = 1\omega_{cp}$ mode of an MCI simulation and the $\omega = 1\omega_{cp}$ mode of a corresponding thermal plasma simulation as a function of time.

and backward propagating waves. The dominant features of the dispersion plot are due to MCI excited waves, that is, the minority proton population entering into cyclotron resonance with the Alfvén wave supported by the bulk plasma. The asymmetry in the dispersion relation must therefore be borne out by the minority ion distribution function. The evolution of this distribution as a function of time is plotted in Fig. 5.14. The data is obtained from the particle momenta as output by the PIC simulations. The upper, middle, and lower rows plot the minority proton distribution function for the x, y, and perpendicular component of momentum respectively. The time at which the distribution functions are computed increases from left to right. In each case, the x-axes are normalised to the maximum value of the relevant component of momenta at that time, while the y-axes are normalised to the maximum value of the distribution function after it has been computed. For instance, in the top left panel, the distribution function of the x-component of the minority proton momenta p_x is calculated at that given time in the simulation. The x-axis is then normalised to the maximum value of p_x , and the y-axis is normalised

to the maximum value of the resulting probability density function before plotting. The magnitude of the perpendicular momentum of the momentum in the equation is calculated using $p_{\perp} = \sqrt{p_x^2 + p_y^2}$, which means the bottom row does not contain any information regarding the direction of motion.

Considering the upper row of Fig. 5.14, we see that the distribution function of the p_x component of the minority proton momentum develops a slight asymmetry as the simulation progresses towards the linearly saturated stage which occurs at around $t \sim 9.5\tau_{gp}$. This asymmetry seems to have mostly disappeared shortly after $t \sim 9.5\tau_{gp}$. The p_y distribution function shows a similar trend, however, the asymmetry is more pronounced just before the linearly saturated stage, and persists for longer into the nonlinear stage of the simulation. Considering the distribution function of p_{\perp} shown in the bottom row of Fig. 5.14, we see that the skew of the distribution function changes as the simulation passes through the point of linear saturation. At late times in the simulation, deep into the nonlinear re-energisation stage, we can see a second bump forming in the distribution function as the minority protons begin to gain energy, a feature that was previously reported in the JET relevant simulations of Ref. [Carbajal et al., 2014].

To further visualise this asymmetry, we plot the x and y components of the bulk velocity of the energetic protons as a function of time. These are shown in the left and right panels of Fig. 5.15 respectively, and are obtained by averaging over all the particle momenta at a given instant in time then dividing by the mass. In both panels the time is normalised to τ_{gp} and the velocity is normalised to the electron thermal velocity $v_{th,e}$. The x-component of the bulk velocity V_x is initially zero, followed by a decrease until $t \sim 10\tau_{gp}$, approximately the point of linear saturation. The y-component of the bulk velocity V_y follows the same trend, however, its bulk velocity increases until $t \sim 10\tau_{gp}$. The decrease in V_x and the increase in V_y are approximately the same. At $t \sim 10\tau_{gp}$, both components sharply reverse their direction, but this change is slightly more noticeable for V_y , which is in line with the features of the distribution functions shown in Fig. 5.14. This asymmetry in the minority energetic proton distribution functions, as well as their bulk velocities, allows for the possibility of backward propagating waves with a different spectral structure to the forward propagating waves. Of course this asymmetry could be purely due to numerical noise and further simulations would have to be done to quantify this, but as mentioned previously, these simulations are extremely well resolved.

We now turn to our attention to bicoherence plots. In Figs. 5.16, 5.17, and 5.18 we plot the cross bicoherence between two signals comprised of the forward

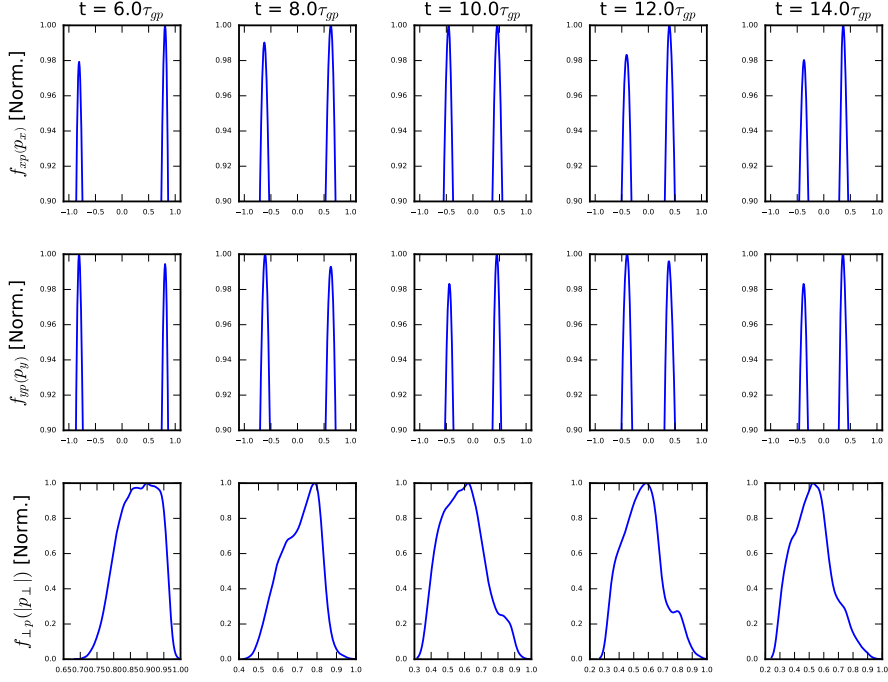


Figure 5.14: Minority proton probability density distribution functions at different times in the simulation. The y-axis of each is normalised to its own maximum probability density and the x-axis each is normalised to its own maximum momentum. Top: p_x probability density function as a function of p_x . Middle: p_y probability density function as a function of p_y . Bottom: $p_{\perp} = \sqrt{p_x^2 + p_y^2}$ probability density function as a function of p_{\perp} .

and backward propagating waves at times $t = 7\tau_{gp}$, $t = 10\tau_{gp}$, and $t = 13\tau_{gp}$ in the simulation. The bicoherence is calculated in wavevector space for a small time-slice in each case. The result is then averaged around this time slice. For example the bicoherence as a function of k is calculated for all times around $t = 7\tau_{gp} \pm 0.25\tau_{gp}$ and is then averaged over this time interval. Nine independent Fourier transforms were used to calculate the bicoherence, giving a significance level of $1/\sqrt{9} \sim 0.33$. Interactions to the right of the vertical white line are between positive modes only. Interactions below the horizontal white line are between negative modes only. Interactions between these lines, in the upper left region of the plot, are between positive and negative modes. If the positive and negative modes were exactly equal in structure, the modes around the diagonal white line would be symmetric, i.e. $|k_1| - |k_2| = |k_2| - |k_1|$. In Fig. 5.16, the instability is still in its




Figure 5.15: Minority energetic proton bulk velocity as a function of time. If the distribution was perfectly symmetric, this would always be zero. Left: Bulk velocity in the x-direction. Right: Bulk velocity in the y-direction. Velocity is normalised to the electron thermal velocity $v_{th,e}$. In both cases we see a distinct change in the bulk velocity to either a net negative or net positive flow at the time of linear saturation (just before $t = 10\tau_{gp}$. This net flow in a preferred direction persists through the non-linear stage, gradually tending to 0.

linear regime and the interactions amongst modes are weak. The dominant nonlinear interactions in the simulation are those between the intense linearly unstable modes in the region of $(k, l) \approx (10, 10)$ and themselves, as well as with a group of modes around $|k| = 25$ (and their backward propagating counterparts). Figure 5.17 plots the bicoherence shortly after linear saturation has been reached at $t = 10\tau_{gp}$. The nonlinear interactions between the modes identified above are much stronger here, and we see a new weaker, but still significant, group of interactions at the forward-backward propagating wave boundary (denoted by the diagonal dashed white line). Figure 5.18 plots the bicoherence well into the nonlinear stage of the simulation at $t = 13\tau_{gp}$. The strongest interactions in the simulations are now less intense, but the interactions at the forward-backward propagating wave boundary are slightly

more intense. The interactions in this region are the most likely cause of the strong nonlinearly driven modes (both forward and backward propagating) that we see at $\omega = 1\omega_{cp}$ in Fig. 5.11.




Figure 5.16: Bicoherence for a time $t = 7\tau_{gp}$ before linear saturation. Strong interactions are not observed and there is almost no asymmetry.

To summarise, we conclude that the peak at $\omega = 1\omega_{cp}$ in the power spectrum shown in Fig. 5.12 is likely caused by nonlinear wave-wave interactions between the strongest linearly unstable modes of forward and backward propagating waves. This mode is linearly stable in the simulation, driven entirely by nonlinear interactions. This could be the origin of the ion cyclotron emission observed at the fundamental proton cyclotron emission frequency as reported in [Ochoukov et al., 2018].




Figure 5.17: Bicoherence for a time $t = 10\tau_{gp}$, shortly after linear saturation. Very strong interactions are observed in several regions, as well as a weaker, but still significant group of interactions at the forward-backward propagating wave boundary.




Figure 5.18: Bicoherence for a time $t = 13\tau_{gp}$, well into the nonlinearly saturated regime of the simulation. The same strong modes as in Fig. 5.17 can be seen, but the strength of the interaction between modes at the forward-backward propagating wave boundary is stronger.

5.3 Conclusions

In this chapter, we first examined the MCI instability in relation to the ICRF ICE JET observed in multiple JET plasma pulses reported in Refs. [Jacquet et al., 2011; McClements et al., 2018]. By running simulations of the MCI in a proton plasma with a minority helium-3 ring beam population we have successfully reproduced the experimentally observed ICE feature at the ^3He fundamental. In addition to this, we examined how different wave propagation angles θ and the inclusion of a finite parallel drift $v_{\parallel 0}$ in the ^3He ring beam distribution affect the simulation dynamics. In the simulations with θ close to 90° , a strong low frequency feature was observed to span a large range of numbers, and deviating away from purely perpendicular wave propagation resulted in this feature being strong enough to resonate with the $l = 3$ ^3He harmonic, thereby increasing its spectral intensity. Other than this, the spectral structure of the simulations remained largely unchanged. The inclusion of a finite $v_{\parallel 0}$ in a simulation with $\theta = 89^\circ$ had little effect on the power spectra, reducing the intensity of some of the harmonics only slightly. It did however, reduce the amplitude of the energy oscillations in the fields and particles, but the overall levels of energy transfer remained approximately the same. We found that for these plasma parameters, increasing θ to 92° led to a predominantly electrostatic instability, and increasing it further beyond 92° meant the instability was no longer excited. We can expect the inclusion of a finite $v_{\parallel 0}$ to be magnified for θ further away from 90° , provided we are still close enough for the MCI to be excited strongly (factoring in computational limitations). This is a potential avenue for future work.

We then analysed the results of a single MCI simulation of the core ICE at the fundamental proton cyclotron frequency in AUG deuterium plasmas [Ochoukov et al., 2018] which may be due to fusion-born protons. The distribution function of freshly born ions in the core is expected to resemble that of a spherical shell. In this simulation, we approximated this as a proton ring beam whose width and position in phase space are initialised using a specific velocity vector $\mathbf{v} = (v_{\perp 0}, v_{\parallel 0})$. Our justification for doing this is that one can think of a spherical shell comprising an infinite number of ring beam distributions each of varying width ($v_{\perp 0}$) and position in the $v_{\parallel 0}$ direction. Only ring beam distributions with $v_{\perp 0}/v_A \sim 1$ are expected to contribute strongly to the MCI resonance. It was found that the fundamental proton cyclotron harmonic was linearly stable in the simulation, but there were many strong MCI excited modes with $l > 5$. Using bicoherence analysis, it was shown that these modes, on both the forward and backward propagating branches of the dispersion relation, couple together nonlinearly to produce many additional spectral

features, including a strong feature at $l = 1$. This $l = 1$ feature was observed to grow in amplitude as the simulation progresses deeper into the nonlinear re-energisation regime, and the strength of the nonlinear coupling was observed to follow a similar pattern. Asymmetries in the components of the minority ion distribution function, as well as the components of the minority proton bulk velocity, lead to asymmetries in the dispersion relation, and hence slight differences in both the linearly and nonlinearly driven modes.

Chapter 6

Simulations of core ICE

The ICE emission considered in the simulations described in chapter 4 was localised to the outer edge region of the KSTAR tokamak, in which drift orbit, local ionisation, and prompt loss effects typically give rise to the highly non-Maxwellian minority ion velocity distributions required for ICE excitation via the MCI. In this region of the plasma, it is appropriate to represent the minority energetic ions as a ring-beam distribution in velocity space [Dendy et al., 1994a,b, 1995]. In this chapter, we turn our attention to ICE emission via the MCI driven by a spherical shell distribution of minority ions [Dendy et al., 1992, 1993], which is appropriate for emission in core tokamak plasmas due to fusion-born ions. Whilst the driving population of the core ICE observed in ASDEX-Upgrade [Ochoukov et al., 2019, 2018], DIII-D [Thome et al., 2018; Heidbrink et al., 2011], and TUMAN-3M [Askinazi et al., 2018] plasma is not completely clear, it seems appropriate to anticipate that some of these ICE signals are driven by fusion-born ions in the core plasma, and examine the feasibility of these ions driving ICE via the MCI.

This chapter is organised as follows: we first describe our simulation set-up, including details of the two non-thermal distribution functions used to represent the minority energetic ions. In Sec. 6.2, we show the results of fourteen independent simulations of the MCI, comparing and contrasting results from ring-beam simulations, typically used in modelling the plasma edge, and spherical shell distributions, typically used in modelling the core plasma. A discussion of the nonlinear aspects of these simulations is presented in Sec. 6.3, with a demonstration of how some minority ion cyclotron harmonics, which are often observed in experiment, can only be explained through nonlinear interactions.

6.1 Description of simulations

We simulate the excitation of ICE by fusion protons using the 1D3V version of the EPOCH PIC code. We simulate fourteen plasmas with parameters closely resembling those used in the calculation of the linear MCI growth rates in Ref. [Dendy et al., 1993], which are relevant to medium-sized tokamak plasmas such as DIII-D and ASDEX-Upgrade. All simulations use an initially uniform electron number density $n_e = 1.0 \times 10^{19} \text{m}^{-3}$, and a background magnetic field $B_z = 3.1 \text{T}$ oriented perpendicular to the simulation domain. The temperatures of the initially Maxwellian background thermal deuterons and electrons are set to 1keV. We denote the bulk deuteron and minority proton number densities by n_D and n_p respectively, and use a fast ion concentration $\xi = n_p/n_D = 10^{-3}$ in all simulations.

In seven of our simulations, the velocity-space distribution of the energetic protons is initialised as a ring-beam distribution according to Eq. 1.26, the parallel velocity and parallel spread have been set to zero, as they are not expected to influence the strictly perpendicularly propagating waves used in this initial study. The value of the initial perpendicular velocity $v_{0\perp}$ corresponds to the 3.02MeV birth speed of fusion-protons produced via D-D reactions. These seven simulation vary the perpendicular spread $v_{T\perp}$ from 0 to $0.3v_{0\perp}$ in steps of $0.05v_{0\perp}$. Note that the $v_{T\perp} = 0.0v_{0\perp}$ distribution function is the same form as that used in previous PIC and hybrid simulations of the MCI [Carbajal et al., 2014; Cook et al., 2013; Chapman et al., 2017, 2018; Carbajal et al., 2017].

Our remaining seven simulations initialise the velocity space distribution of energetic protons according to a spherical shell distribution according to Eq. 1.27. In a similar fashion to the ring-beam simulations, the initial particle speed v_0 corresponds to the 3.02MeV birth speed of fusion-protons, and we vary the velocity spread v_T (shell thickness) from 0 to $0.3v_0$ in steps of $0.05v_0$. Note that most of the figures contained within this chapter do not display the results of simulations with $v_T = 0.05v_0$ and $v_{T\perp} = 0.05v_{0\perp}$, as they are similar in nature to simulations with $v_T = 0.1v_0$ and $v_{T\perp} = 0.1v_{0\perp}$. They are however shown in figures in which contour plots are used, ensuring uniform interpolation between results from different simulations. Each of the seven ring-beam simulations lasts 40 proton gyro-periods τ_{cp} , by which time the instability is well into its nonlinear saturated regime. The seven spherical shell simulations last for varying lengths of time, from $70\tau_{cp}$ to $130\tau_{cp}$, depending on the time it takes the instability to reach linear saturation.

6.2 Linear physics

Let us first examine the variation in the energies of the fields and particles in our simulations. Fig. 6.1 displays the change in energy density as a function of time for six MCI simulations using a ring-beam velocity distribution for the minority protons. The perpendicular velocity spread $v_{T\perp}$ expressed as a fraction of the initial perpendicular velocity $v_{0\perp}$ is shown at the top of each panel. The energy transfer between particles and fields qualitatively resembles that of previous work [Carbajal et al., 2014; Cook et al., 2013], with the minority protons transferring their energy to the bulk plasma and to the fields. However, the timescale of the instability (normalised to the minority proton gyro-period $\tau_{cp} = 2\pi/\omega_{cp}$) is almost four times longer than anything previously reported. This can be explained in terms of analytical growth rates of the MCI for these simulation parameters, which are discussed later in the section and are of the order $\sim 10^{-2}\omega_{cp}$, much lower than that of previous simulations [Carbajal et al., 2014; Cook et al., 2013]. The simulation with no perpendicular velocity spread $v_{T\perp}$ reaches linear saturation the earliest, at around $25\tau_{cp}$, and displays the lowest level of collective oscillation between the bulk plasma and the fields. As $v_{T\perp}$ is increased, this oscillation becomes larger in amplitude, and the linear phase of the instability saturates at a later time. Once $v_{T\perp}$ is increased beyond $0.1v_{0\perp}$, the amount of energy transferred from the minority proton population to the fields and bulk plasma drops off rapidly and affects both the field components in the same way, with the peak changes in E_x and B_z energy densities being approximately five times less in the $v_{T\perp} = 0.3v_{0\perp}$ simulation than they are in the $v_{T\perp} = 0.1v_{0\perp}$ simulation. In all panels we see a nonlinear re-energisation stage shortly after linear saturation which persists for the remainder of the simulation. This stage is least pronounced for the $v_{T\perp} = 0.0v_{0\perp}$ simulation, and most pronounced for the $v_{T\perp} = 0.1v_{0\perp}$ simulation, lasting for a smaller fraction of the total simulation time as $v_{T\perp}$ increases further.

The corresponding six plots for simulations in which the minority protons were initialised according to a spherical shell distribution are shown in Fig. 6.2. The top three panels are qualitatively the same as their ring-beam counterparts, providing the first direct simulation evidence that ICE can be excited via the MCI when the minority ions are represented by a spherical shell distribution. All three of these panels show a distinct linear saturation phase, followed by re-energisation, the amount of which decreases substantially as the velocity spread v_T increases. The amount of energy transfer shown in all panels of Fig. 6.2 is more than ten times

less than that of their ring-beam counterparts, and in all cases the field amplitude saturation time is at least double that of the ring-beam simulations. The bottom three panels differ from the top three in that they do not enter a re-energisation stage. Instead, the magnetic field amplitude saturates at about $t = 80\tau_{gp}$ in all cases, and the minority protons continue transferring the remainder of their energy solely to the bulk plasma deuterons, for what appears to be an indefinite amount of time.




Figure 6.1: Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time from six PIC simulations in which the minority energetic protons are initialised using ring-beam velocity distributions with varying perpendicular velocity spreads. The perpendicular spread $v_{T\perp}$ as a fraction of the initial perpendicular velocity $v_{0\perp}$ is shown at the top of each panel. The traces, ordered from top to bottom at their peak (and in colour online) are: Top (red) the change in kinetic energy density of the thermal bulk plasma deuterons; second (green) the energy density of the magnetic field perturbation ΔB_z ; third (blue) the energy density of the electrostatic field E_x ; fourth (cyan) the change in kinetic energy density of the minority energetic protons. Time is normalised to the proton gyro period. Note the different scales on the y-axis of each plot.




Figure 6.2: Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time from six PIC simulations in which the minority energetic protons are initialised using spherical shell velocity distributions with varying velocity spreads (shell thickness). The spread v_T as a fraction of the initial velocity v_0 is shown at the top of each panel. The ordering of the traces and the normalisation is the same as in Fig. 6.1

The advantage of the long duration of both sets of simulations is that we are able to obtain exceptional frequency resolution when performing Fourier transforms, and are able to see clearly, without any ambiguity, that in all cases the MCI is excited under these conditions. Figure 6.3 shows the distribution of energy in the z -component of the magnetic field in frequency-wavenumber space, where the Fourier transform has been performed over the entire duration of the simulation and the full simulation domain. Each panel is plotted using the same \log_{10} scale, and the maximum displayed frequency is approximately equal to the lower hybrid frequency $\omega_{LH} \approx 9\omega_{cp}$, where ω_{cp} is the proton cyclotron frequency. Each panel shows a series of well defined resonances along the magnetoacoustic branch at integer multiples of ω_{cp} . These resonances are strongest for the simulations with little or no perpendicular spread, which is expected as Fig. 6.1 shows these simulations

have the greatest energy transfer from the minority protons to the magnetic field. In addition to the resonances along the magnetoacoustic branch, we see regions of high spectral density at $\omega = 7\omega_{cp}$ and $\omega = 8\omega_{cp}$ at lower wavenumber. The intensity of these regions decreases as $v_{T\perp}$ increases. These additional modes are nonlinear in origin, and arise because of intense phase coupling between modes on the main magnetoacoustic branch. We will revisit this later in this section, but for now it is important to note that these modes have similar intensity to their linearly excited counterparts, so any spectrally structured ICE at these high frequencies is due at least in part to strong nonlinear interactions. The power spectrum for these simulations is shown in Fig. 6.4. We see a strong mode at $\omega = 4\omega_{cp}$ in all panels except for the one with $v_{T\perp} = 0.3v_{0\perp}$. At low values of $v_{T\perp}$, we see only a very weak $\omega = 5\omega_{cp}$ mode, which becomes more powerful as $v_{T\perp}$ increases, in conjunction with a decrease in the amplitude of a mode at $\omega = 4\omega_{cp}$.

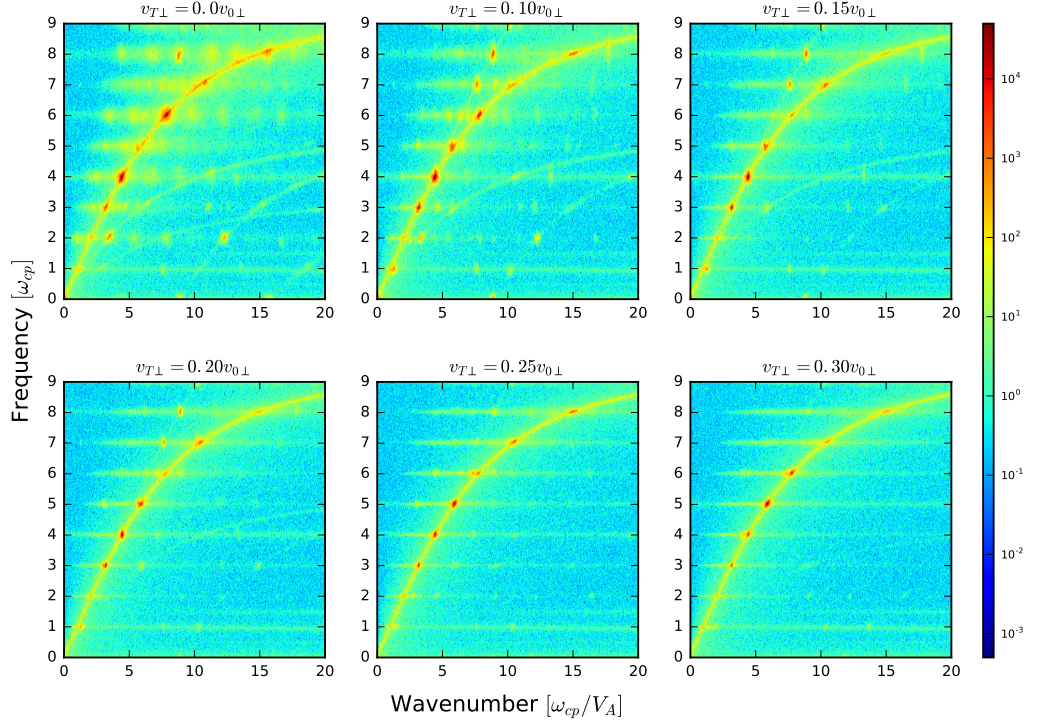


Figure 6.3: Distribution of energy in the fluctuating z-component of the magnetic field ΔB_z across frequency-wavenumber space from six PIC simulations in which the minority energetic protons are initialised using ring-beam velocity distributions with varying perpendicular velocity spreads. The spread $v_{T\perp}$ as a fraction of the initial perpendicular velocity $v_{0\perp}$ is shown at the top of each panel. The plots are a spatio-temporal Fourier transform of the B_z field component and in each case span the entire spatial domain and the length of time shown in their Fig. 6.1 counterparts. Shading indicates the \log_{10} of the spectral density of the oscillatory part ΔB_z of the B_z field component. The vertical axes are normalised to the minority proton cyclotron frequency ω_{cp} , while the horizontal axes are normalised to ω_{cp} divided by the Alfvén speed V_A .

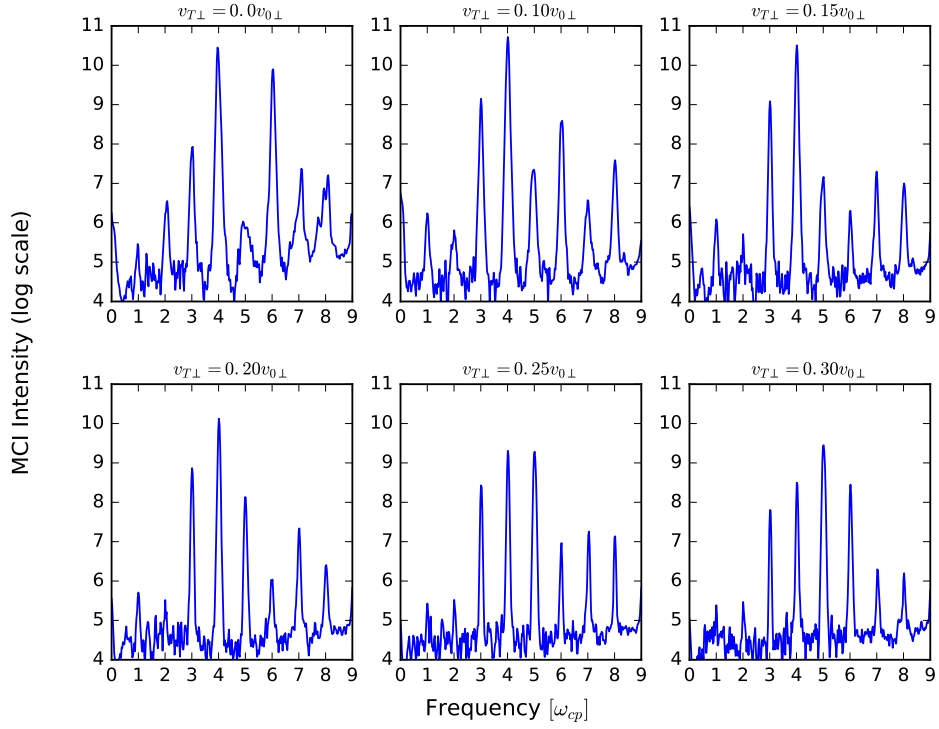


Figure 6.4: Spectral intensity of the fluctuating B_z field energy density, from six PIC simulations in which the minority energetic protons are initialised using ring-beam velocity distributions with varying perpendicular velocity spreads. The spread $v_{T\perp}$ as a fraction of the initial perpendicular velocity $v_{0\perp}$ is shown at the top of each panel. The vertical axes are plotted on a \log_{10} scale while the horizontal axes are normalised to the minority proton cyclotron frequency ω_{cp} . The range of both axes is identical in all panels.

The corresponding dispersion relation and power spectrum plots for the spherical shell distribution simulations are shown in Figs. 6.5 and 6.6 respectively. The frequency resolution of these plots is much greater than that of their ring beam counterparts, owing to a longer simulation time. As such it is perhaps difficult for one to pick out the strong, highly localised resonances in the dispersion relation. We can however see that in the $v_T = 0.0v_0$ simulation, there is a faint mode at $\omega = 8\omega_{cp}$ to the left of the main dispersion branch, just as in its ring-beam counterpart. It is much more informative to look at the power spectrum shown in Fig. 6.6, in which the spectral structure is slightly different from that of the ring-beam simulations. As before the $\omega = 1\omega_{cp}$ and $\omega = 2\omega_{cp}$ modes are not present, and the $\omega = 4\omega_{cp}$ is by far the most dominant for the the simulations with low or zero v_T . There is no gap at the $\omega = 5\omega_{cp}$ mode in these simulations, and by the time v_T has been increased to $0.2v_0$, all modes are roughly equal in magnitude. As v_T increases further, the dominant mode shifts to the $\omega = 6\omega_{cp}$ mode, much like the dominant mode in the ring-beam simulations shifts upwards to $\omega = 5\omega_{cp}$. Unlike the ring-beam simulations, the $\omega = 7\omega_{cp}$ and $\omega = 8\omega_{cp}$ modes are more similar in magnitude to the spectrally intense modes at lower frequencies.

We now turn our attention to the linear growth rates of modes in both sets of simulations, which are plotted as a function of cyclotron harmonic number and thermal spread in the right panels of Fig. 6.7. The left panels display the analytical linear growth rates computed numerically using a first principles kinetic dispersion solver [Irvine, 2018, 2019]. The upper panels of Fig. 6.7 correspond to a spherical shell distribution, and the growth rates are plotted as a function of the shell spread v_T , while the lower panels correspond to a ring-beam distribution, and are plotted as a function of the perpendicular velocity spread $v_{T\perp}$. To obtain uniform spacing of the contours two additional PIC simulations were run, corresponding to $v_T = 0.05v_0$ in the case of a spherical shell distribution and $v_{T\perp} = 0.05v_0$ in the case of a ring-beam distribution. These two simulations are not shown in the other figures of this paper, and do not alter any conclusions drawn. In all panels the growth rates are normalised to the proton cyclotron frequency ω_{cp} , and the eight harmonic is not displayed as both its analytical and PIC linear growth rates are either zero or tiny in all cases - this mode is almost entirely driven by nonlinear interactions, as we will show later.




Figure 6.5: Distribution of energy in the fluctuating z -component of the magnetic field ΔB_z across frequency-wavenumber space from six PIC simulations in which the minority energetic protons are initialised using spherical shell velocity distributions with varying velocity spreads (shell thickness). The spread v_T as a fraction of the initial velocity v_0 is shown at the top of each panel. The plots are a spatio-temporal Fourier transform of the B_z field and in each case span the entire spatial domain and the length of time shown in their Fig. 6.2 counterparts. All other details are the same as Fig. 6.3.




Figure 6.6: Spectral intensity of the fluctuating B_z field energy density, from six PIC simulations in which the minority energetic protons are initialised using spherical shell velocity distributions with varying velocity spreads (shell thickness). The spread v_T as a fraction of the initial velocity v_0 is shown at the top of each panel. All other details are the same as Fig. 6.4.

The ring-beam growth rates in the simulations are qualitatively identical to their analytical counterparts; strikingly, the absence of the fifth harmonic in PIC simulations with low $v_{T\perp}$ is accurately predicted by the analytical growth rate solver. The growth of the sixth harmonic at low values of $v_{T\perp}$ is also consistent, as is the tendency for the power to shift from the fourth harmonic to the fifth as $v_{T\perp}$ increases. There is a discrepancy in the magnitude of the growth rates, with the maximum value in the simulation being ~ 1.7 times larger than the maximum analytical value. This is due to the overwhelmingly large growth rate of the fourth harmonic in the simulation, which is not predicted by the analytical solver. The spherical shell distribution growth rates in the PIC simulations are both qualitatively and quantitatively consistent with their analytical counterparts. Both analytical theory and simulation predict that the growth is concentrated around the fourth proton cyclotron harmonic at low values of v_T , and shifts towards the sixth harmonic as v_T increases. The numerical solver predicts that the fourth and fifth modes have growth rates which are similar in magnitude, which is not borne out by the simulation results. It also predicts that the largest growth rate at $v_T = 0.3v_0$ occurs at the fourth harmonic, while it is the sixth harmonic in the simulations. These discrepancies between analytical and simulation growth rates could be due to slightly different wave propagation angles. As discussed in section 6.1, all our simulation are initialised with a propagation angle of $\theta = 90^\circ$, whereas the numerical solver calculates growth rates at $\theta = 89.9^\circ$, owing to the pole at $\theta = 90^\circ$ which must be avoided. Even with these discrepancies, the similarity between analytical and simulation growth rates across fourteen simulations with two different energetic proton velocity distributions is excellent.

The linear growth rates for the spherical shell distribution have also been calculated according to Eq. 31 of [Dendy et al., 1993] and are shown in Fig. 6.8. This analytical expression predicts the same trend in the magnitude of the growth rate as a function of v_T , however the peak growth rate is located at the sixth harmonic as opposed to the fourth harmonic in the simulations and analytical solver. We note that the tendency of the fastest growing mode in the simulations to shift from the fourth, to the fifth, and finally to the sixth proton cyclotron harmonic as the velocity spread increases is qualitatively the same in both the ring-beam and spherical shell simulations. The only major difference in the simulation growth rates is the presence of faster growing $\omega = 6\omega_{cp}$ modes in the ring-beam simulations.




Figure 6.7: Contour plot displaying the linear growth rates of the MCI as a function of velocity spread and proton cyclotron harmonic number. Left panels: Linear growth rates are calculated numerically from a kinetic dispersion solver. Right panels: Linear growth rates calculated directly from PIC simulations. Upper panels: Linear growth rates corresponding to a minority proton ring-beam distribution function with varying perpendicular velocity spread $v_{T\perp}$. Lower panels: Linear growth rates corresponding to a minority proton spherical shell distribution function with varying velocity spread v_T . The colour bar above each panel shows the magnitude of the growth rate γ normalised to the proton cyclotron frequency ω_{cp} . In all panels the eighth proton cyclotron harmonic is omitted, as it is found to have a zero or very small linear growth rate in all cases. In the upper panels, corresponding to the ring-beam distribution, the numerical and PIC growth rates differ slightly in their maximum and minimum values, but the left and right plots are qualitatively almost identical. In the lower panels, corresponding to the spherical shell distribution, the numerical and PIC growth rates have the same maximum and minimum values, but the left and right plots are not as qualitatively similar as their ring-beam counterparts.

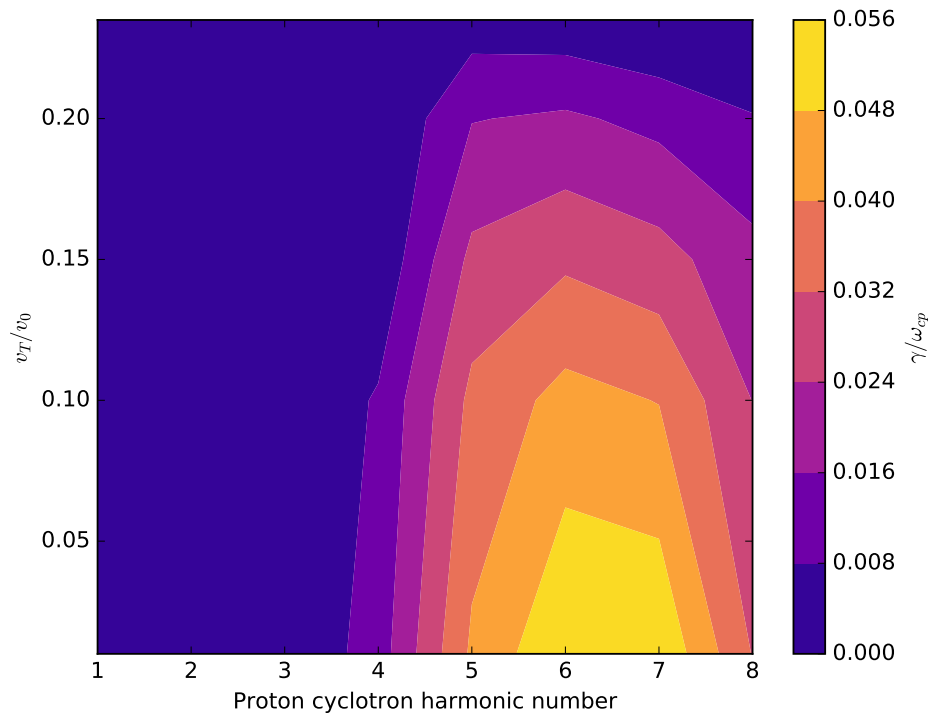


Figure 6.8: Contour plot displaying the linear growth rates of the MCI as a function of velocity spread v_T and proton cyclotron harmonic number. The minority protons are represented using a spherical shell distribution and the growth rate, calculated according to Eq. 31 of [Dendy et al., 1993], is normalised to the proton cyclotron frequency ω_{cp} . This analytically calculated growth rate is maximum around the sixth proton cyclotron harmonic and decreases as the shell thickness increases.

6.3 Nonlinear physics

We now discuss the nonlinear aspects of the simulations. The evolution of the fluctuating part of the z -component of the magnetic field ΔB_z as a function of time and wavenumber is shown in Fig. 6.9 for two simulations of the MCI. The left panel of Fig. 6.9 is from a ring-beam simulation with $v_{T\perp} = 0.0v_0$, while the right panel of Fig. 6.9 is from a spherical shell simulation with $v_T = 0.0v_0$. In both panels time is normalised to τ_{cp} , wavevector is normalised to ω_{cp}/V_A , and shading indicates the \log_{10} of the spectral density of ΔB_z . In both panels we see the strongest modes around $k \sim 4.35\omega_{cp}/V_A$ and $k \sim 7.8\omega_{cp}/V_A$, which begin growing in the linear stage of the simulation and correspond to the fourth and sixth proton cyclotron harmonics respectively. In addition, we see a strong, linearly driven mode at $k \sim 5.9\omega_{cp}/V_A$ in the right panel, corresponding to the fifth cyclotron harmonic. This mode also appears in the left panel, but weaker and at a later time, in the nonlinear phase of the simulation. In both panels we see a relatively strong nonlinearly driven mode at $k \sim 8.7\omega_{cp}/V_A$ mode, which contributes to the eighth harmonic in both simulations, along with a $k \sim 15.5\omega_{cp}/V_A$ mode, which is entirely nonlinearly driven in the left panel, and is linearly driven but incredibly weak in the right panel. Other modes occurring at late times can be seen growing in both panels, mostly the left panel. The evolution of B_z^2 as a function of time for the $k \sim 8.7\omega_{cp}/V_A$ and $k \sim 15.5\omega_{cp}/V_A$ modes in the $v_{T\perp} = 0.0v_0$ ring-beam simulation is shown in Fig. 6.10. The traces were obtained by averaging over a narrow range of wavevectors in the vicinity of these modes, being careful not to have any overlap with any other, clearly distinct modes shown in Fig. 6.9. A moving average was also applied in time to smooth out the field oscillations. The $k \approx 8.7\omega_{cp}/V_A$ mode reaches a slightly higher saturation energy and grows slightly later in time than the $k \approx 15.5\omega_{cp}/V_A$ mode. The upper left panel of Fig. 6.3 shows that both these modes have a strong resonance at $\omega = 8\omega_{cp}$, along the eighth proton cyclotron harmonic, and thus both contribute in comparable magnitude to the $\omega = 8\omega_{cp}$ peak shown in the top left panel of Fig. 6.4.

We now quantify the extent of this nonlinear wave-wave coupling in the simulations. The best quantitative evidence for this is obtained from bispectral analysis, which is described in detail in Sec. 2.2 of chapter 2. Figure 6.11 plots the square of the bicoherence b^2 for six MCI simulations in which the minority protons were initialised according to a ring-beam distribution. The squared bicoherence is plotted as a function of wavenumber, which, for the strongest modes, maps directly to frequency. Shading indicates the intrinsic strength of nonlinear coupling, 1 (dark




Figure 6.9: Distribution of energy in the fluctuating z -component of the magnetic field ΔB_z across time-wavenumber space. Left: From a PIC simulation in which the minority energetic protons are initialised using a ring-beam velocity distribution with no perpendicular velocity spread. Right: From a PIC simulation in which the minority energetic protons are initialised using a spherical shell velocity distribution with no velocity spread. Both plots show the spatial Fourier transform of the B_z field and span the entire spatial domain. Shading indicates the \log_{10} of the spectral density of the oscillatory part ΔB_z of the B_z field component. The vertical axes are normalised to the minority proton cyclotron frequency ω_{cp} while the horizontal axes are normalised to ω_{cp} divided by the Alfvén speed V_A .

red) being completely coupled and 0 (dark blue) completely uncoupled. In each plot b^2 was calculated using the full 40 proton gyro periods τ_{cp} of simulation data, the width of each successive Fourier transform was $1.25\tau_{cp}$, and the overlap of each successive Fourier transform was $0.25\tau_{cp}$. The number of independent samples is thus 32, giving a significance level of $b^2 > 1/\sqrt{32} \gtrsim 0.177$, which is far lower than the coupling strength of the modes of interest. For simplicity, in the following we refer to frequencies by their proton cyclotron harmonic number l , e.g. $\omega = 4\omega_{cp}$ becomes $l = 4$. We will also assume that it is implicit that the wavenumbers are normalised, e.g. “ $k = 4.35$ ” actually refers to $k = 4.35\omega_{cp}/V_A$. We refer to parent modes using the subscripts 1 and 2, and child modes using the subscript 3, as in Sec. 2.2. Note that a child mode (k_3, l_3) can subsequently interact with other modes and become the parent to new child modes, in which case we then refer to the old child modes, now parent modes, using the subscripts 1 and 2.




Figure 6.10: Time evolution of the $k \approx 8.7\omega_{cp}/V_A$ and $k \approx 15.5\omega_{cp}/V_A$ field components of a PIC simulation in which the minority energetic protons are initialised using a ring-beam distribution with no perpendicular velocity spread, corresponding to the left panel of Fig. 6.9. A moving average has been applied in time in order to smooth out the field oscillations (see the green trace in Fig. 6.1). Both modes begin to gain energy close to the start of the nonlinear stage of the simulation, meaning they are likely driven by nonlinear interactions.

All six panels of Fig. 6.11 display multiple regions of strong nonlinear wave-wave coupling, and, in general, the strength of this coupling decreases as $v_{T\perp}$ increases, perhaps in part due to the shorter duration of the nonlinear re-energisation stage of the simulation. In all panels a mode at $(k_1, l_1) \approx (4.35, 4)$ interacts with all other modes strongly, and the strength of this interaction decreases rapidly with increasing $v_{T\perp}$. This mode interacts strongly with itself, and is the origin of the $(k_3, l_3) \approx (8.7, 8)$ mode. This child mode is also phase coupled with all other modes, and can in theory produce “3rd generation” child modes. The strength of coupling between the $(k, l) \approx (8.7, 8)$ and other modes drops off rapidly as $v_{T\perp}$ increases, following the behaviour of its parent. The strength of the phase coupling between $(5.9, 5)$ and other modes (the vertical line from $k_2 = 5.9$ and the horizontal line from $k_2 = 5.9$) is weak at low $v_{T\perp}$, and greatly increases as $v_{T\perp}$ increases and this mode becomes linearly unstable. The same can be said of the $(10.5, 7)$ mode (see Figs. 6.7 and 6.4).

We can further characterise the nature of nonlinear interactions in these six ring-beam simulations by examining the summed bicoherence Σb^2 . This is plotted in Fig. 6.12 for each ring-beam simulation. The vertical axis sums all the squared bicoherence contributions from modes k_1 and k_2 which add vectorially to produce modes with $k_3 = k_1 + k_2$. If a wave $k = k_3$ has a large value of Σb^2 , then there are, in general, multiple combinations of k_1 and k_2 with large values of b^2 that contribute to the formation of k_3 . Note this does not necessarily mean a wave with $k = k_3$ is observed in the simulation, it could be the case that all its parents have a low amplitude so that in practice, the child mode is never observed. One must also remember that any combinations that satisfy $k_3 = k_1 + k_2$ must also satisfy $l_3 = l_1 \pm l_2$, where the positive and negative signs refer to forward and backward propagating waves respectively. Combinations of k_1 and k_2 such that $b^2 < 0.8$ are not included in the sum, meaning only modes with the strongest nonlinear couplings are considered.

Considering, for example, the $v_{T\perp} = 0.0v_0$ panel of Fig. 6.12, we see spikes at $k_3 \approx 8.7, 12.2, 13.1, 15.5,$ and 18.7 . The $k_3 \approx 8.7$ spike is the nonlinearly driven mode which contributes to the $l = 8$ proton cyclotron harmonic and is present in all but the $v_{T\perp} = 0.3v_{0\perp}$ panel. The large spike at $k_3 \approx 12.2$ corresponds to a mode that can be clearly seen in the left panel of Fig. 6.9. Its parent modes can be identified as $(k_1, l_1) \approx (7.8, 6)$ and $(k_2, l_2) \approx (4.35, 4)$. Interestingly, this $k_3 \approx 12.2$ mode corresponds to two different waves, one at $l_3 = l_1 + l_2 = 6 + 4 = 10$, and $l_3 = l_1 - l_2 = 6 - 4 = 2$, the latter can be clearly seen in the top left panel of Fig.

6.3. We will return to the $(k_3, l_3) \approx (12.2, 10)$ mode presently. Another noticeable spike occurs at $k_3 \approx 15.6$. Perhaps the most intuitively obvious combination of k_1 and k_2 that contributes to this spike is a self interaction between the strong linearly stable $(7.8, 6)$ mode. To satisfy energy conservation this would have to manifest itself as an $(k_3, l_3) \approx (15.6, 12)$ mode. This is indeed the case, and is shown in the top left panel of Fig. 6.13 along with the $(k_3, l_3) \approx (12.2, 10)$ discussed earlier, and other modes that are the result of nonlinear phase coupling. These modes are above the lower hybrid frequency $\omega_{LH} \approx 9\omega_{cp}$, a region where linearly excited modes are evanescent [Chapman et al., 2018; Cairns, 1985; Verdon et al., 2009]. The same high frequency regions in ring-beam simulations with finite $v_{T\perp}$ show similar nonlinearly driven modes, whose existence can be explained in terms of the corresponding finite $v_{T\perp}$ panels of Fig. 6.12. The number of these high frequency non-linearly excited modes decreases as $v_{T\perp}$ increases, reflecting the decreasing strength of nonlinear coupling. A plot of the MCI intensity as a function of frequency for this high frequency region is shown in Fig. 6.14. Considering the top left panel, only the $l = 10$ and $l = 12$ are above the noise level of ~ 4 orders of magnitude, which is deduced from Fig. 6.4. The $l = 10$ mode is over two orders of magnitude above the noise level, and is of comparable magnitude to the linearly unstable mode at $l = 7$ in the same simulation. This $l = 10$ mode persists in the simulations with finite $v_{T\perp}$, albeit to a lesser degree.

While nonlinear self interaction between the strong linearly stable $(7.8, 6)$ mode certainly contributes to the $k_3 \approx 15.6$ spike in the top left panel of Fig. 6.12, it does not explain the $(k_3, l_3) \approx (15.6, 8)$ mode which is nonlinearly driven in the $v_{T\perp} = 0.0v_0$ simulation (it is very weakly linearly unstable in the finite $v_{T\perp}$ simulations). It is difficult to find a combination of k_1 and k_2 which sum to ~ 15.6 , whilst satisfying the frequency matching requirement. A candidate pair of modes is the two nonlinearly driven modes $(k_1, l_1) \approx (12.2, 10)$ and $(k_2, l_2) \approx (3.4, 2)$, such that $l_3 = 10 - 2 = 8$. The $(k_2, l_2) \approx (3.4, 2)$ mode can be seen in the top left panel of Fig. 6.3, and in wavevector space is just about distinguishable from the linearly unstable $(k, l) \approx (3.3, 3)$ mode - if we look closely in this region of the left panel of Fig. 6.9, we can see a mode evolving during the nonlinear stage of the simulation which is directly adjacent to the linearly stable mode at $k \approx 3.3$. This is a clear example of two nonlinearly driven modes interacting with each other to produce another nonlinearly driven mode which is in the range of frequencies below ω_{LH} , which are often detectable in experiment. The spike at $k_3 \approx 13.1$ can also be explained in terms of secondary interactions between nonlinearly driven modes. This

is due to the nonlinearly excited $k_1 \approx 8.7$ mode interacting with the linearly unstable $k_2 \approx 4.35$ mode; the latter being the parent of the former. This manifests itself as a weak but clearly visible $(k, l) \approx (13.1, 12)$ mode in the upper left panel of Fig. 6.12.




Figure 6.11: The square of the bicoherence b^2 of the oscillatory part of the B_z field component as a function of normalised wavenumber k from six PIC simulations in which the minority energetic protons are initialised using ring-beam velocity distributions with varying perpendicular velocity spreads. The spread $v_{T\perp}$ as a fraction of the initial perpendicular velocity $v_{0\perp}$ is shown at the top of each panel. The colour scale indicates intrinsic nonlinear coupling between waves with wave numbers k_1 and k_2 , which takes values between 0 and 1. In each plot b^2 was calculated from the full 40 proton gyro periods τ_{cp} of simulation data, the width of each successive Fourier transform was $1.25\tau_{cp}$, and the overlap of each successive Fourier transform was $0.25\tau_{cp}$. The number of independent samples is thus 32, giving a significance level of $b^2 > 1/\sqrt{32} \gtrsim 0.177$.




Figure 6.12: The sum of the square of the bicoherence b^2 as a function of normalised wavenumber k_3 from six PIC simulations in which the minority energetic protons are initialised using ring-beam velocity distributions with varying perpendicular velocity spreads. The spread $v_{T\perp}$ as a fraction of the initial perpendicular velocity $v_{0\perp}$ is shown at the top of each panel. The sum is calculated from corresponding data shown in Fig. 6.11, and $k_3 = k_1 + k_2$ is the child mode resulting from strong linear coupling between two parent modes k_1 and k_2 . Only combinations of k_1 and k_2 yielding a value of $b^2 \geq 0.8$ (see Fig. 6.11) are included in the sum. A large value of Σb^2 indicates that a child mode in the vicinity of k_3 arises from strong nonlinear coupling between multiple distinct parent modes k_1 and k_2 .




Figure 6.13: As Fig. 6.3, except only the region with frequency ω greater than the lower hybrid frequency ω_{LH} is shown. In this region we see the child modes with wavenumber $k_3 = k_1 + k_2$ and frequency $\omega_3 = \omega_1 + \omega_2$ that arise from the strong nonlinear wave-wave interactions shown in Figs. 6.11 and 6.12.




Figure 6.14: As Fig. 6.4, except only the region with frequency ω greater than the lower hybrid frequency ω_{LH} is shown. This plot is obtained by integrating over the region shown in Fig. 6.13. In each panel, not all the peaks are uniform in height, suggesting that some of the more intense spectral peaks, e.g. $\omega = 10\omega_{cp}$ in the top left panel, owe their existence to the strong nonlinear wave-wave interactions shown in Figs. 6.11 and 6.12. This mode in particular, which is driven entirely by nonlinear wave-wave interactions, has a spectral intensity of comparable magnitude to some of the linearly driven modes shown in Fig. 6.4. Thus this mode and others like it are the direct result of the minority energetic proton population relaxing under the MCI, and not merely normal modes of the system.

The squared bicoherence for six MCI simulations in which the minority protons were initialised according to a spherical shell distribution is shown in Fig. 6.15 and their nonlinear characteristics closely resemble their ring-beam counterparts. Unlike the ring-beam simulations, the duration of the spherical shell simulations is highly variable, so the bicoherence shown in each panel of Fig. 6.15 has to be computed using time series of different length. In each panel, b^2 was calculated using time series spanning the range of data displayed in the corresponding panels shown in Fig. 6.2. For instance, the bicoherence shown in the top left panel is calculated using 73 proton gyroperiods of data, the bicoherence shown in the top middle is calculated using 80 proton gyroperiods of data, and so on. The minimum significance level is thus $b^2 \gtrsim 0.13$.

In the top row of simulations, the $k_3 \approx 8.7$ mode interacts strongly with all modes, again owing to its strong parent mode at $(k, l) \approx (4.35, 4)$. This coupling is barely evident in the $v_T = 0.2v_0$ simulation, and is non-existent in the two $v_T > 0.2v_0$ simulations. This is because the parent $(k, l) \approx (4.35, 4)$ is not strongly coupled to anything in these simulations, owing to a characteristically different nonlinear stage, see panels $v_T = 0.25v_0$ and $v_T = 0.3v_0$ of Fig. 6.2. It is immediately obvious that the $v_T = 0.0v_0$ simulation has the weakest (but still statistically significant) coupling of all the spherical shell simulations. Despite this, the nonlinearly excited $(k, l) \approx (8.7, 8)$ mode is still strongest in the $v_T = 0.0v_0$ simulation, evidenced by the top left panel of Fig. 6.5. This is because its parent, the $(k, l) \approx (4.35, 4)$ mode is most powerful in this simulation, almost an order of magnitude more powerful than the simulations with finite v_T . It is important to remember that even weak phase coupling between linearly excited modes of high amplitude can produce noticeably strong nonlinearly driven modes.

As v_T increases, the $(k, l) \approx (5.8, 5)$ mode becomes more strongly coupled to other modes in the simulation, mirroring the evolution of its amplitude with v_T . This is almost exactly the behaviour of the ring-beam simulations, as the velocity spread increases, the mode grows, as does the strength of its coupling to other modes. This shift towards high l modes continues as v_T is increased further, again exhibiting similar behaviour to the ring-beam simulations. As before, we plot the sum of the squared bicoherence as a function of wavenumber for these spherical shell simulations. This is shown in Fig. 6.16, in which only combinations of k_1 and k_2 yielding a value of $b^2 \geq 0.6$ are included in the sum for all simulations except the $v_T = 0.0v_0$. This simulation exhibits significantly weaker coupling than its finite v_T

counterparts, so instead values with $b^2 \geq 0.4$ are included in the sum. The range of wavevectors plotted in Fig. 6.16 is less than that of Fig. 6.12, owing to a more narrow range of wavevectors that satisfy the criteria detailed above. One must also bear in mind that the scale of the y-axis is shorter than in Fig. 6.12.




Figure 6.15: The square of the bicoherence b^2 of the oscillatory part of the B_z field component as a function of normalised wavenumber k from six PIC simulations in which the minority energetic protons are initialised using spherical shell velocity distributions with varying velocity spreads. The spread v_T as a fraction of the initial velocity v_0 is shown at the top of each panel. In each panel b^2 was calculated using time series spanning the range of data displayed in the corresponding panels shown in Fig. 6.2. For instance, the bicoherence shown in the top left panel is calculated using 73 proton gyroperiods of data, the bicoherence shown in the top middle is calculated using 80 proton gyroperiods of data, and so on. The minimum significance level is thus $b^2 \gtrsim 0.13$. We note that the strength of coupling in the simulation with $v_T = 0.0v_0$ shown in the top left panel is significantly less than the rest of the simulations.




Figure 6.16: The sum of the square of the bicoherence b^2 as a function of normalised wavenumber k_3 from six PIC simulations in which the minority energetic protons are initialised using spherical shell velocity distributions with varying velocity spreads. The spread v_T as a fraction of the initial perpendicular velocity v_0 is shown at the top of each panel. The sum is calculated from corresponding data shown in Fig. 6.15, and $k_3 = k_1 + k_2$ is the child mode resulting from strong linear coupling between two parent modes k_1 and k_2 . For the upper left panel, which corresponds to a simulation with $v_T = 0.0v_0$, only combinations of k_1 and k_2 yielding a value of $b^2 \geq 0.4$ (see Fig. 6.15) are included in the sum. For the rest of the panels, that is the simulations with finite velocity spread, only combinations of k_1 and k_2 yielding a value of $b^2 \geq 0.6$ are included in the sum.

As expected, we see spikes at $k_3 \approx 8.7$ in simulations with $v_T < 0.25v_0$, much like the ring-beam simulations. We also see spikes at $k_3 \approx 12.2$, 13.1, 15.5, and 18.7 across a wide range of simulations. These spikes are also seen in the ring-beam simulations, albeit with stronger coupling. In the $v_T = 0.0v_0$ simulation, we see the same example of nonlinearly driven modes interacting to produce other modes as in the corresponding ring-beam simulation, that is, $k_3 \approx 13.1 = k_2 + k_1 \approx 4.35 + 8.7$. Unlike the ring-beam simulation, no strong modes above the noise level are observed above the lower hybrid frequency, likely because the intensity of the parent modes

is significantly less than that of their ring-beam counterparts. It is possible that a simulation with more computational macro-particles could improve upon this result by decreasing the noise level, but at present this would require a tremendous amount of computing resources.

6.4 Conclusions

In this chapter we have performed the first PIC simulations of the MCI due to the collective relaxation of minority energetic protons modelled using an isotropic spherical shell distribution. These first principles PIC simulations self-consistently solve the Maxwell-Lorentz equations for fully kinetic thermal ion, energetic ion, and electron populations, and are allowed to progress deep into the nonlinear regime. This choice of distribution function is significantly different from the delta-function minority ion distributions used in the works of [McClements et al., 2018; Carbajal et al., 2014; Cook et al., 2013; Carbajal et al., 2017; Cook et al., 2017; Chapman et al., 2017, 2018], and is a possible model distribution function for fusion-born ions in the core of tokamak plasmas [Dendy et al., 1992, 1993; D. J., 1979]. We have analysed six simulations of the MCI under these conditions, varying the velocity spread in each one, as well as six additional simulations of the MCI in which the minority protons were initialised using ring-beam distribution functions of varying perpendicular velocity spread.

We found that the MCI was excited in all cases, and the spherical shell simulations took around twice as long to saturate as their ring beam counterparts, with around ten times less energy transfer from the minority ions to the bulk plasma and electromagnetic fields. In general, the time to linear saturation was found to increase as the velocity spread increased in both the spherical shell and ring-beam simulations. However, at thermal spreads $v_T \geq 0.2v_0$, the spherical shell simulations exhibited a fixed linear saturation time of approximately 80 proton gyro-periods, with the saturation energy decreasing as v_T increased, and the remainder of the free energy in the minority ion distribution being transferred approximately monotonically to the bulk deuterons.

The mode structure in both sets of simulations was found to be qualitatively similar, with the intensity of ICE in the ring-beam simulations sometimes two orders of magnitude more than that of the spherical shell simulations, owing to the increased energy transfer between the minority protons and the fields. In each case,

increasing the velocity spread led to the most spectrally intense mode gradually shifting from the fourth proton cyclotron harmonic, to the fifth, and finally to the sixth. This suggests that by observing the mode structure in experiments, one may be able to deduce the spread of the minority ion distribution, an important parameter as it determines how much energy is transferred to the bulk plasma and electromagnetic fields due to the MCI. The linear growth rates of modes in all twelve simulations were found to closely resemble the analytical linear growth rates, which were calculated using a first principles kinetic dispersion solver [Irvine, 2018, 2019].

The nonlinear aspects of each simulation were discussed, and many significant nonlinear wave-wave couplings were identified. In general, there are many more wave-wave interactions in the ring-beam simulations than in the spherical shell simulations, however, the strongest nonlinearly driven modes of practical interest, i.e. those that contribute significantly to the ICE signal, are present among both sets of simulations. In particular, both sets of simulations exhibit a strong nonlinearly driven $(k, l) \approx (8.7, 8)$ mode which does not lie along the magnetoacoustic dispersion branch, and in the case of the ring-beam simulation with zero velocity spread, contributes to approximately half of the total intensity of the eighth proton cyclotron harmonic. Other nonlinear couplings gave rise to modes above the lower hybrid frequency ω_{LH} , and in the ring-beam simulations, some of these modes had intensities comparable to low intensity linearly unstable modes. In the $v_{T\perp} = 0.0v_{0\perp}$ ring-beam simulation, it was suggested that one of these modes at $(k, l) = (12.2, 10)$ couples to another nonlinearly driven mode at $(k, l) \approx (3.3, 2)$, providing a second source of energy to the eighth proton cyclotron harmonic, which is thus entirely driven by nonlinear interactions. This demonstrates how indispensable the nonlinear physics is when simulating ICE and interpreting experimental observations. The key to identifying mode couplings was to first fulfil the wavenumber matching criterion, followed by the frequency matching criterion. These two requirements highlight that a modest experimental effort to detect both the perpendicular wavenumber and high frequency ion cyclotron harmonics would enable us to better understand the measured frequency spectrum, and hence the character of the energetic ion distribution function.

The similarity between: the variation of energy density with spread, the linearly excited mode structure, and the nonlinear characteristics of both sets of simulations, suggests that a ring-beam velocity distribution for the minority energetic ions serves as a close approximation to an isotropic spherical shell distribution,

provided the velocity spread isn't too large. From a resource perspective this is crucial. In this paper, both sets of simulations used large numbers of computational macro particles and ran for a long time, resulting in high quality spectra. Firstly, the ring-beam simulations are already around half as computationally demanding as their spherical shell counter parts, as they take half as long to reach linear saturation. Secondly, the intensity of ICE harmonics in the ring-beam simulations is much higher than that of the spherical shell simulations, meaning in future simulations, significantly less computational particles could be used whilst still maintaining a satisfactory signal-to-noise ratio. The diagnostics with which to measure core ICE in tokamaks are becoming more widespread, and the computing resources with which to simulate it are becoming increasingly more sophisticated; we are thus at a juncture in which it will soon be feasible for PIC simulations of the MCI to be used for predictive modelling of tokamak plasma phenomena, as opposed to only interpretive modelling. The "cheaper" ring-beam simulations offer a way to realise this.

Chapter 7

Simulations with multiple minority ion species

In this chapter we discuss PIC simulations of the MCI in which there are two minority ion species. In section 7.1 we analyse the results of two simulations relevant to the KSTAR tokamak that have both a minority fusion-born proton population, and a minority neutral beam injected (NBI) deuteron population. The purpose of this is to provide a link between the fusion-born proton ICE discussed in Sec. 4.1 and Sec. 4.3, and the deuteron NBI ICE discussed in Sec. 4.4.

In section 7.2, we discuss preliminary simulations of the MCI in JET tokamak plasmas in which both a minority NBI deuteron population and a minority low temperature helium ash population are present. These simulations describe a mechanism that might enable us to pump out the alpha particle ash in deuterium-tritium tokamak plasmas, which can be detrimental to plasma confinement [Wesson, 2004].

7.1 Simulations of KSTAR ICE with NBI deuterons and fusion-born protons

Chapter 4 addresses two distinct types of ICE that are observed in KSTAR plasmas:

1. Steady-state ICE in steps of the deuteron cyclotron frequency preceding the pedestal collapse through multiple ELM filament bursts, discussed in Sec. 4.1 and Sec. 4.3.
2. Highly dynamic downward frequency chirping ICE in steps of the proton cyclotron frequency during the pedestal collapse, discussed in Sec. 4.4.

The MCI simulations in chapter 4 which pertain to these two phenomena, consider them to be distinct, and model each type of ICE with two separate sets of MCI simulations. Each set of simulations uses a single minority ion species represented by a ring-beam distribution function: NBI deuterons in the case of 1), and fusion-born protons in the case of 2). This is valid because these two types of ICE are seldom observed in the same KSTAR pulse. On the rare occasions when they are, the steady-state deuteron ICE occurs $\sim 50\mu s$ before the chirping proton ICE, which is approximately 100 times longer than the duration of a typical PIC simulation of the MCI. It seems appropriate to conjecture that these two types of ICE events naturally follow each other during ELM crashes, at least in some of the KSTAR pulses. The steady-state NBI ICE exists before the ELM crash and can be reproduced with PIC simulations using a background electron number density similar to that found at the top of the pedestal. As the pedestal collapses through multiple filament bursts, the local density changes, so the ICE is no longer steady-state and instead chirps down, which can be reproduced by running PIC simulations with a range of electron number densities. It is therefore possible that both types of ICE exist simultaneously at the moment before the pedestal collapse, which we refer to as $t = t_0 = 0$ following the convention of chapter 4. The purpose of this section is to test this hypothesis by simulating these two types of ICE simultaneously, and to ascertain to what extent one effect influences the other, making sure that any conclusions are consistent with our previous results in chapter 4.

We run a PIC simulation of the MCI using one spatial dimension and three velocity dimensions. Both energetic deuterons and energetic protons are present in our simulation, and are represented by ring-beam distribution functions. The magnetic field $B_{0z} = 1.44\text{T}$ is oriented along the z-axis, perpendicular to our simulation domain, and the background plasma density $n_e = 2.5 \times 10^{19}\text{m}^{-3}$ reflects that found at the top of the pedestal in a typical KSTAR plasma. We set the background electron and thermal deuteron temperatures $T_e = T_D = 1\text{keV}$, the ratio of NBI deuterons to bulk plasma ions $\xi_{NBI} = 10^{-3}$, and the ratio of fusion-born protons to bulk plasma ions $\xi_p = 10^{-3}$. The NBI deuterons are initialised as a 100keV ring-beam with initial pitch angle of 72.4° , zero drift, and no thermal spread, and are therefore identical to the NBI deuteron population used in the simulations described in Sec. 4.4. The fusion-born protons are initialised as a 300keV ring-beam with zero drift and no thermal spread, and are therefore identical to the fusion-born proton population used in the simulations of Sec. 4.1 and Sec. 4.3. The NBI deuterons are allowed to evolve independently of the fusion-born protons for the first

10 deuteron gyro-periods τ_{cD} of the simulation. At $t = 10\tau_{cD}$, the NBI deuterons have reached both linear and nonlinear saturation, and the fusion-born protons are “switched on” and evolved in tandem with the NBI deuterons for a further $8\tau_{cD}$. The simulation uses 50500 grid cells, and 800 particles per cell, meaning that for the first 10 deuteron gyro-periods, the simulation is identical to the $B_{0z} = 1.44\text{T}$ simulation discussed in Sec. 4.4, the only difference being this new simulation has less computational macro-particles. At $t = 10\tau_{cD}$, when the protons are initialised, the simulation set-up very closely resembles that which makes up the left-most strip of the lower left panel of Fig. 4.3 of Sec. 4.1. The key difference is that instead of the fusion-born protons being initialised in a thermal plasma as in Sec. 4.1, they are initialised in a plasma in which the wave and particle processes due to a minority NBI deuteron population relaxing under the MCI have already been established.

Figure 7.1 shows the change in energy density as a function of time, normalised to τ_{cD} , for this simulation. At early times, before the protons are initialised, the primary energy flow from the NBI deuterons is to the thermal deuterons, whose kinetic oscillation helps support the field oscillations excited by the MCI. These field oscillations include, with comparable magnitude, an electromagnetic component $(\Delta B_z)^2$, and an electrostatic component E_x^2 . This stage of the instability reaches linear saturation by $t \approx 2\tau_{cD}$ and nonlinear saturation by $t \approx 7\tau_{cD}$. At $t = 10\tau_{cD}$, the minority fusion-born protons are initialised, and they begin to “re-excite” the MCI in what appears to be a two-stage process. From $t = 10\tau_{cD}$ to $t \approx 13\tau_{cD}$, the fusion-born protons transfer their energy to the fields and bulk plasma in a manner similar to the NBI deuterons, characteristic of ICE excitation via the MCI. In addition, some of this energy is transferred to the now dormant NBI deuterons. From $t \approx 13\tau_{cD}$ until the end of the simulation at $t \approx 18.5\tau_{cD}$, there is a second excitation stage, in which the magnetic field component B_z and the bulk deuterons gain energy, further exciting the MCI. Unlike the first stage of this process, there is no more energy transfer to the E_x field component. The NBI deuterons continue to be energised by the fusion-born protons, resulting in significantly less field excitation than was present in the purely fusion-born proton driven ICE discussed in Sec. 4.1.




Figure 7.1: Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time, from a PIC simulation with $\xi_{NBI} = 10^{-3}$ and $\xi_p = 10^{-3}$. The traces, ordered from top to bottom at their first peak ($t \approx 2\tau_{cD}$), are: Top (red) the change in kinetic energy density of the thermal bulk plasma deuterons; second (green) the energy density of the magnetic field perturbation ΔB_z ; third (blue) the energy density of the electrostatic field E_x ; fourth (cyan) the change in kinetic energy density of the minority energetic fusion-born protons, which are not initialised until $t = 10\tau_{cD}$; fifth (magenta) the change in kinetic energy density of the minority energetic NBI deuterons. Time is normalised to the deuteron gyro period.

The spatiotemporal Fourier transform of this simulation using data from the first 10 deuteron gyro-periods (20 proton gyro-periods τ_{cp}) is shown in Fig. 7.2. The y-axis is normalised to the deuteron cyclotron frequency ω_{cD} ($= 0.5 \times \omega_{cp}$), and the x-axis is normalised to ω_{cD} divided by the Alfvén speed v_A . All other frequencies and wavenumbers in the rest of this section use the same normalisation. Figure 7.2 shows no clear ICE excitation and the dispersion relation resembles that of a thermal plasma. This was the case in Sec. 4.4, in which a higher value of ξ_{NBI} was required to raise the signal above the noise level.




Figure 7.2: Distribution of energy in the fluctuating z -component of the magnetic field ΔB_z across frequency-wavenumber space from a PIC simulation with $\xi_{NBI} = 10^{-3}$ and $\xi_p = 10^{-3}$. This plot is a spatio-temporal Fourier transform of the B_z field over the intervals spanning $0 \leq x \leq 50500\lambda_D$ and $0 \leq t \leq 10\tau_{cD}$ ($20\tau_{cp}$), before the fusion-born protons have been initialised. Shading indicates the \log_{10} of the spectral density of the oscillatory part ΔB_z of the B_z field component in frequency-wavenumber space. It is apparent that no visible ICE excitation is present, due to the high levels of noise (see Sec. 4.4). As such, this dispersion relation resembles that of a thermal plasma.

In Fig. 7.3, we re-plot the dispersion relation using data from the entire duration of the simulation, $18.5\tau_{cD}$ ($37\tau_{cp}$). The corresponding power spectrum, obtained by integrating over wavenumber, is shown in Fig. 7.4. We can see multiple spectral peaks at integer multiples of ω_{cp} (the even integer multiples of ω_{cD}), including the $\omega = 17\omega_{cp}$ peak that is present in the left-most strip of the similar simulation shown in the lower left panel of Fig. 4.3.




Figure 7.3: As Fig. 7.2, except the temporal part of the Fourier transform of the B_z field is over the interval spanning the full range of the simulation $0 \leq t \leq 18.5\tau_{cD}$ ($37\tau_{cp}$). Spectrally intense regions at harmonics of the proton cyclotron frequency (even harmonics of the deuteron cyclotron frequency) are present.

To try overcome the simulation noise present in the first $10\tau_{cD}$ of the simulation, before the protons are initialised, we now run a second simulation, increasing both minority ion fractions to $\xi_{NBI} = 10^{-2}$ and $\xi_p = 10^{-2}$. As in Sec. 4.4, this is not expected to affect our conclusions, because the simulated ICE power due to the MCI has been found to scale linearly with fast particle concentration [Carbajal et al., 2017]. The energy dynamics for this simulation are shown in the left panel of Fig. 7.5. This plot looks rather peculiar because of huge amount of energy transfer between the minority protons and minority deuterons at late times in the simulations. A close up of this plot, focussing on the magnetic field excitation, is shown in the right panel of the same figure. The first $10\tau_{cD}$ are similar to that of Fig. 7.1, however, the instability is much more rapid due to the large value of $\xi_{NBI} = 10^{-2}$, yielding a larger linear growth rate than that of the simulation with $\xi_{NBI} = 10^{-3}$. Despite the rapid saturation of the instability, this NBI deuteron-only portion of the simulation ran for $10\tau_{cD}$, allowing us to obtain the frequency resolution necessary to




Figure 7.4: The spectral intensity of the fluctuating B_z field energy density obtained by integrating Fig. 7.3 over wavenumber. Strong spectral peaks at harmonics of the proton cyclotron frequency (even harmonics of the deuteron cyclotron frequency) are present.

distinguish the closely spaced deuteron cyclotron harmonics. As before, the protons are initialised at $t = 10\tau_{cD}$ and excite the z-component of the magnetic field, the thermal deuterons, but transfer the majority of their energy to the NBI deuterons, mediated by the magnetic field. The overall energy transfer to the magnetic field at $t > 10\tau_{cD}$ is again much greater than that at $t < 10\tau_{cD}$, and also takes place in two stages, up to $t = 15\tau_{cD}$. Beyond $t = 15\tau_{cD}$, the simulation has reached linear saturation, and the energy oscillates between the magnetic field, the bulk deuterons, and the energetic deuterons for the remainder of the simulation.

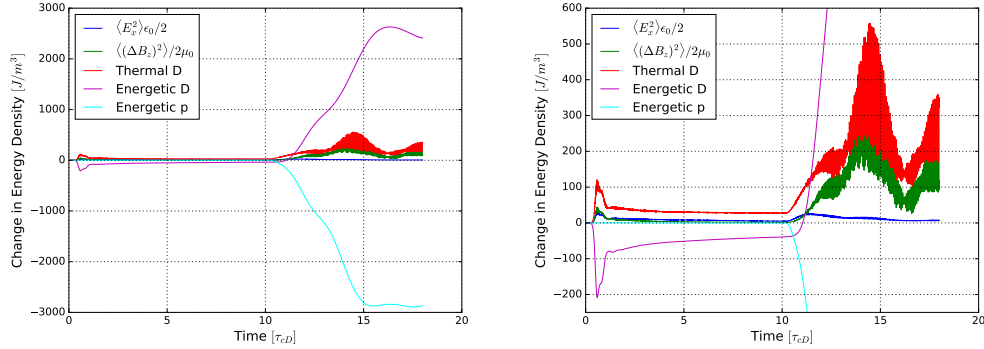


Figure 7.5: Left: As Fig. 7.1 but from a simulation with $\xi_{NBI} = 10^{-2}$ and $\xi_p = 10^{-2}$. Right: An inset of the left panel, focussing on the evolution of the z-component of the magnetic field B_z .

We know from panel c) of Fig. 4.22, that ICE at deuteron cyclotron harmonics is observed under these conditions, but for completeness, the power spectrum as a function of frequency corresponding to the first $10\tau_{cD}$ ($20\tau_{cp}$) from our new simulation is shown in Fig. 7.6. Deuteron cyclotron harmonics above $\omega = 14\omega_{cD}$ are lost in the simulation noise and hence are not shown here. We can confidently say that despite the lower number of computational macro-particles in comparison to the simulation whose results are shown in Fig. 4.22, we observe ICE with spacing ω_{cD} in our new simulation. As in Sec. 4.1, it becomes difficult to distinguish the high frequency harmonics from the background noise.

The power spectrum using data from the full $18\tau_{cD}$ ($36\tau_{cp}$) of the simulation is shown in Fig. 7.7. The deuteron cyclotron harmonics are difficult to distinguish at low frequencies, and at high frequencies, which are experimentally relevant [Kim et al., 2018; Thatipamula et al., 2016], the spectrum is completely dominated by ICE with spacing equal to the proton cyclotron frequency. The spectral peaks at multiples of the proton cyclotron frequency are up to three orders of magnitude greater than ICE at harmonics of the deuteron cyclotron frequency, and in some cases, slightly shifted from integer multiples of ω_{cp} . The cause of this shift is unknown, and would require more simulations and analysis to understand, but it could be due to nonlinear wave-wave interactions between deuteron and proton cyclotron harmonics, or simply the result of doing a Fourier transform on what is essentially a composite, highly non-stationary signal. These questions are left for future studies.

These tentative results may suggest that in a situation in which both energetic NBI deuterons and fusion-born protons are present, both types of ICE are also present, but fusion-born proton ICE is by far the most spectrally intense, so




Figure 7.6: The spectral intensity of the fluctuating B_z field energy density from a simulation with $\xi_{NBI} = 10^{-2}$ and $\xi_p = 10^{-2}$. Power is obtained by performing a spatio-temporal Fourier transform of the B_z field over the intervals spanning $0 \leq x \leq 50500\lambda_D$ and $0 \leq t \leq 10\tau_{cD}$ ($20\tau_{cp}$) and then integrating over wavenumber. The temporal range of simulation data corresponds to times before the initialisation of the minority fusion-born protons.

much so that NBI deuteron ICE is barely observed. These results may therefore explain the apparent absence of ICE with spacing ω_{cD} in KSTAR spectrograms that exhibit downward frequency chirping with spacing ω_{cp} . These results may also explain the apparent side-band features present at high frequencies in the upper panel of Fig. 4.4 of Sec. 4.1 (repeated here as Fig. 7.8 for convenience), although without many more simulations of higher quality than those presented in this section, it is impossible to draw any reasonable conclusions as to the extent of this effect.

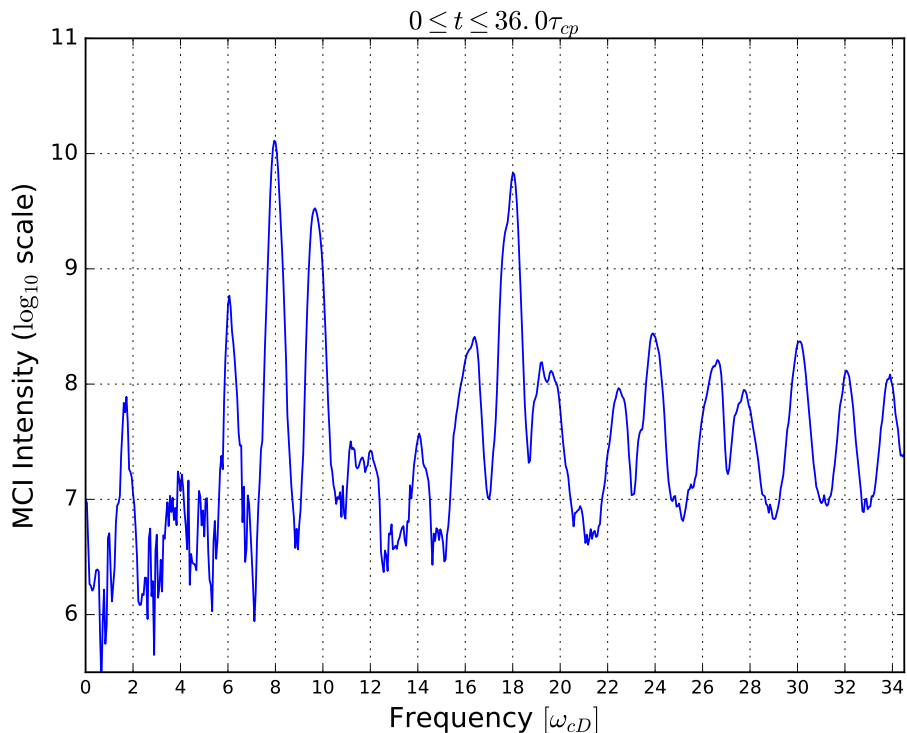


Figure 7.7: As the Left panel of Fig. 7.6, except the temporal part of the Fourier transform of the B_z field component is over the interval spanning the full range of the simulation, $0 \leq t \leq 18\tau_{cD}$ ($36\tau_{cp}$), after the fusion-born protons have been initialised. ICE at harmonics of the proton cyclotron frequency (even harmonics of the deuterion cyclotron frequency) is more intense than ICE at harmonics of the deuterion cyclotron frequency.

7.2 Preliminary simulations of helium ash pumping in JET core plasmas

In future deuterium-tritium plasmas, such as those in ITER, some of the α -particles produced in fusion reactions may accumulate in the core region in the form of helium ash. The presence of this ash is highly undesirable for two main reasons; one, it increases radiation losses via bremsstrahlung radiation due to the increased acceleration of electrons in the presence of ions with high charge numbers [Wesson, 2004]; and two, it dilutes the plasma fuel, limiting the maximum achievable plasma density and hence the fusion power. The removal of helium ash is therefore of tremendous practical interest for future burning plasmas. There are a variety of mechanisms under consideration for the removal of helium ash, including transport




Figure 7.8: Experimentally-measured fast RF burst spectrogram from KSTAR plasma 11513 with $B_0 = 1.99\text{T}$ and average electron number density before the ELM crash $\langle n_e \rangle = 2.6 \times 10^{19}\text{m}^{-3}$. The white dashed lines denote successive proton cyclotron harmonics. Downward step-wise frequency chirping with $\omega_{cp}/2\pi = f_{cp} \sim 25\text{MHz}$ is apparent, and some of the features are accompanied by side-band features of lesser intensity.

from the core to the edge via Sawtooth oscillations [Chapman, 2010], pumping via ion cyclotron resonance heating (ICRH) [Hamamatsu et al., 1998], and using resonant magnetic perturbations (RMPs) to control the helium entry into the core plasma at the edge [Schmitz et al., 2016]. In addition, ELMs in H-mode plasmas may offer a natural way to expel impurities, such as helium ash from the exhaust region [Kamiya et al., 2007; Zohm, 1996].

In this section, we discuss preliminary simulation results which may pave the way for another method with which to pump out the helium ash. The basic idea is to use high energy perpendicular NBI to “kick” helium ash particles in the core plasma onto a new trajectory with a larger perpendicular velocity component. In a tokamak plasma, helium ash particles with $v_{\perp} \gg v_{\parallel}$ may then lie on orbits which take them outside of the plasma and into the exhaust region. We run a series of 1D3V PIC simulations using plasma parameters typical of JET core plasmas. The temperature of the thermal electrons and deuterons is set to $T_e = T_D = 5\text{keV}$, the background plasma density $n_e = 9.8 \times 10^{19}\text{m}^{-3}$, and the magnetic field $B_z = 2.7\text{T}$ is oriented along the z-axis, perpendicular to the simulation domain. Each simulation uses 10,150 grid cells with 1,000 particles per cell. We represent our minority NBI deuterons using a ring-beam velocity distribution with zero parallel velocity and zero thermal spread. Helium ash is represented using a Maxwellian velocity distribution. We set the ratio of helium ash to bulk ions $\xi_{\alpha} = 10^{-3}$, and the ratio of NBI deuterons to bulk ions $\xi_{NBI} = 10^{-3}$.

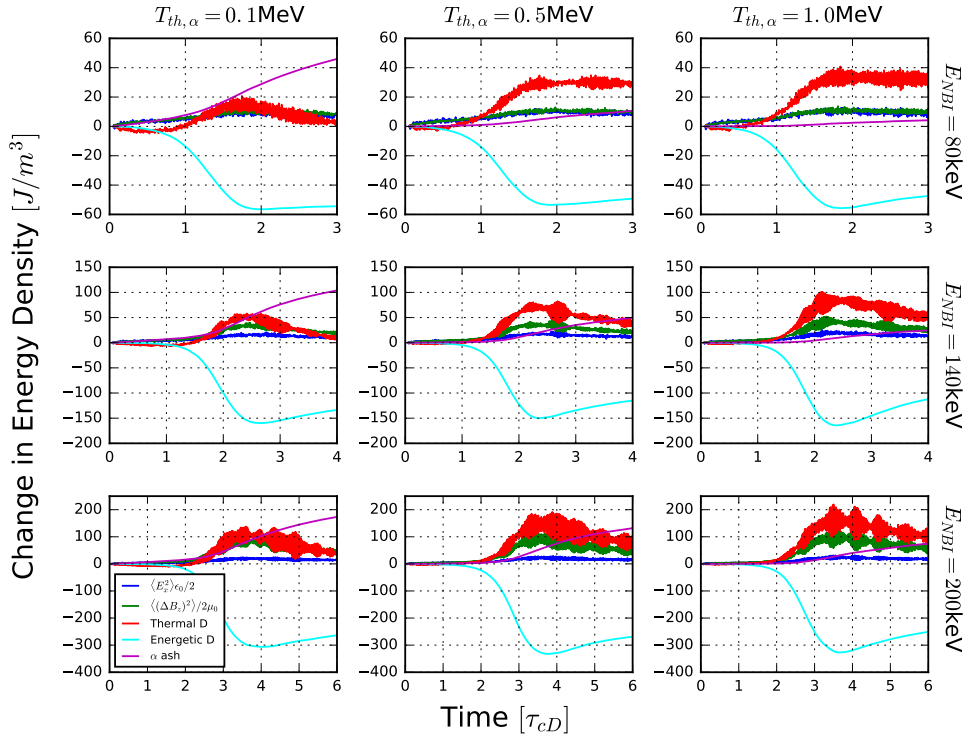


Figure 7.9: Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time in multiple PIC simulations with initial NBI deuteron energies 80keV, 140keV, and 200keV (rows), and initial helium ash temperatures 0.1MeV, 0.5MeV, and 1.0MeV (columns). The traces, ordered from top to bottom in the upper left panel are: Top (magenta) the change in kinetic energy density of the minority helium ash; second (red) the change in kinetic energy density of the thermal bulk plasma deuterons; third (green) the energy density of the magnetic field perturbation ΔB_z ; fourth (blue) the energy density of the electrostatic field E_x ; fifth (cyan) the change in kinetic energy density of the minority energetic NBI deuterons. Time is normalised to the deuteron gyro period.

A total of nine simulations were run using a combination of three initial NBI deuteron energies: $E_{NBI} = 80\text{keV}$, 140keV , and 200keV ; and three values for the temperature of the helium ash: $T_{th,\alpha} = 0.1\text{MeV}$, 0.5MeV , and 1.0MeV . The change in energy density as a function of time for these nine simulations is shown in Fig. 7.9, in which rows correspond to a fixed value of E_{NBI} and share the same range of x and y axis values, and columns correspond to a fixed value of $T_{th,\alpha}$. In all simulations, the helium energy density, denoted by the magenta trace, increases as time progresses. This additional energy comes directly from the NBI deuterons, whose velocity is entirely in the perpendicular direction, and the energy transfer is

mediated by the bulk ions and the electromagnetic fields. The amount of energy transfer to the helium ash increases as the NBI deuteron energy increases, as does the time for the simulation to reach linear saturation. Crucially, the influence of the NBI deuterons is greatest for low temperature helium ash - those which contribute most to the dilution of the plasma and the reduction in the maximum achievable fusion power. The most extreme example of this is the simulation with $E_{NBI} = 80\text{keV}$ and $T_{th,\alpha} = 1.0\text{MeV}$ shown in the top left panel of Fig. 7.9, in which the simulation largely behaves as if the helium ash were absent.

For completeness, Fig. 7.10 shows the results of a simulation with $E_{NBI} = 140\text{keV}$ and $T_{th,\alpha} = 0.1\text{MeV}$, but with a helium ash concentration $\xi_\alpha = 10^{-6}$, a thousand times less than the deuterium NBI concentration. Under these conditions, there is no energy transfer to the helium ash, the energy dynamics resemble those typical of the MCI.




Figure 7.10: Time evolution of the change in energy density of particles and electric and magnetic fields as a function of time in a PIC simulation in which the initial NBI deuteron energy is 140keV, the initial helium ash temperature is 0.1MeV, $\xi_{NBI} = 10^{-3}$, and $\xi_\alpha = 10^{-6}$. Time is normalised to the deuteron gyro period. No appreciable energy transfer takes place between the NBI deuterons and the helium ash.

We now turn our attention to the perpendicular velocity distribution functions of the helium ash in the nine simulations which make up the panels of Fig. 7.11. These distribution functions are calculated from particle data at the time shown in the top right corner of each panel, which is approximately equal to the time of linear saturation in each simulation. The x-axes have been normalised to the ensemble average of the magnitude of the initial perpendicular velocity of the NBI deuterons. For the simulations with $T_{th,\alpha} = 0.1\text{MeV}$, shown in the first column, there is a distinct bump in the distribution function at $v_{\perp,\alpha} \approx v_{\perp NBI}$, meaning at the time of linear saturation, the momentum transfer is localised in velocity space. This bump in the velocity distribution is present in the bottom two panels of the middle column, but is much less pronounced. Thus, only NBI deuterons with high energy are able to affect significant change in a helium ash population which has not yet fully cooled. There are no bumps present in the right column corresponding to $T_{th,\alpha} = 1.0\text{MeV}$, meaning the most energetic helium particles are left unaffected by this process.

7.3 Conclusions

In this chapter we have carried out simulations of the MCI relevant to two novel scenarios in which there are two minority ion species present. The first of these pertains to ICE observed during KSTAR plasmas, where the pedestal is observed to collapse through multiple filament bursts during ELM crashes. We ran two simulations with different minority ion concentrations, initialising both minority NBI deuterons and minority fusion-born protons using ring-beam velocity distributions, the latter being time delayed with respect to the former. We found that in both simulations the fusion-born proton ICE was much more spectrally intense than the NBI deuteron ICE, offering a potential explanation as to why steady-state ICE in steps of the deuteron cyclotron frequency, and highly dynamic downward frequency chirping ICE in steps of the proton cyclotron frequency, are not observed at the same time during the KSTAR ELM crash process. We note in passing that given enough computational resources, it may be possible to show that faint deuteron ICE at high frequencies is observed alongside the dominant proton ICE, which could provide an explanation for existence of the side-band features observed in some KSTAR plasmas that exhibit downward frequency chirping in steps of the proton cyclotron frequency.

We then discussed preliminary results of PIC simulations pertaining to helium ash pumping in JET core plasmas. We identified a novel mechanism in which

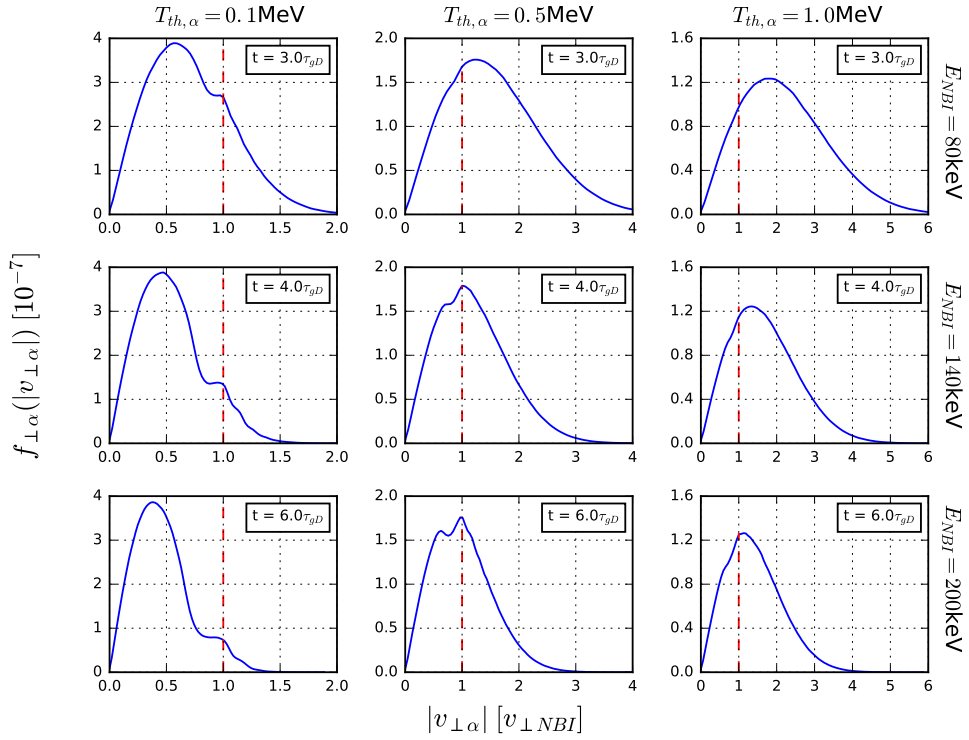


Figure 7.11: Perpendicular velocity distribution function of the helium ash population in multiple PIC simulations with initial NBI deuteron energies 80keV, 140keV, and 200keV (rows), and initial helium ash temperatures 0.1MeV, 0.5MeV, and 1.0MeV (columns). The y-axes are normalised to 10^{-7} , while the x-axes are normalised to the ensemble average of the magnitude of the initial velocity of the NBI deuteron population. In each panel, the distribution function was calculated using particle data outputted directly from the PIC simulation at the time of linear saturation, which corresponds to the troughs of the cyan traces shown in Fig. 7.9, and is displayed in the top right inset of each panel.

perpendicular deuteron NBI can be used to selectively target helium ash in the core, transferring a considerable amount of perpendicular momentum to a subset of the ash population, which may alter its orbit in way which carries them out of the core plasma. For an initial NBI deuteron energy of 140keV, the highest energy that can be realistically used to heat JET plasmas, only low temperature helium ash is targeted by this mechanism, leaving the still mildly energetic helium unaffected. This means that it may be possible to modify existing NBI systems to pump out helium ash in the core, and it may even be worthwhile to have a dedicated perpendicular NBI system whose sole purpose is for helium ash pumping, and not necessarily plasma heating. This effect is also observed at a comparatively low NBI injection

energy of 80keV, meaning it may be possible to observe using the current NBI systems in place on medium sized tokamaks such as KSTAR or ASDEX-Upgrade. The simulations presented in this section are the subject of ongoing research, and further work is underway to quantify the effect of NBI beams on helium ash populations in core tokamak plasmas.

Chapter 8

Summary

In this thesis we have used particle in cell (PIC) simulations for the numerical study of Ion cyclotron emission (ICE) excited by the magnetoacoustic cyclotron instability (MCI) in the context of three tokamak devices. In chapter 4 we investigated two distinct types of ICE in the KSTAR tokamak, fusion product (FP) driven ICE at spacing of the proton cyclotron frequency, and ICE with spacing equal to the deuteron cyclotron frequency driven by neutral beam injection (NBI) [Thatipamula et al., 2016; Kim et al., 2018]; and two preliminary simulations attempting to establish a synergy between these two types of ICE were analysed in chapter 7. In chapter 5, we carried out simulations of ICE in the JET tokamak, which was excited by a minority ^3He population heated by ICRF waves. We then performed a single simulation of the MCI with ASDEX Upgrade (AUG) plasma parameters to ascertain if the MCI could explain the FP ICE with spacing equal to the proton cyclotron frequency [Ochoukov et al., 2018], concluding that it is indeed a likely emission mechanism. Motivated by the contemporary core ICE results [Ochoukov et al., 2018, 2019], in chapter 6 we carried out a numerical study of MCI excitation due to both ring beam and spherical shell distributions of varying thicknesses, providing the first precedent for ICE excitation due to the MCI from spherical shell distributions, which arise in core plasmas. I will now offer brief summaries of the main findings in each chapter.

8.1 KSTAR ICE

During Edge Localised Mode (ELM) crashes in KSTAR deuterium plasmas, bursts of spectrally structured ICE are detected. We first examined KSTAR chirping ICE with spacing equal to the proton cyclotron frequency. This frequency chirping was

observed on sub-microsecond timescales near the plasma edge. By first using orbit calculations, we were able to prove that a subset of fusion born protons in KSTAR remain on confined, deeply passing orbits. We then carried out multiple simulation runs for different, adjacent, values of the plasma density under KSTAR edge conditions, which enabled us to infer the theoretical dependence of ICE spectral structure on the local electron number density. By matching this density dependence to the observed time-dependence of chirping ICE spectra in KSTAR, we obtained sub-microsecond time resolution of the evolving local electron number density during the ELM crash.

It was clear from this analysis that the lower hybrid frequency ω_{LH} was playing a substantial role in the simulation dynamics. To this end, we undertook a numerical study to quantify to what extent ω_{LH} influenced the dynamics of the simulation and experiment. We found that, for several values of magnetic strength, decreasing the density below a certain critical value resulted in a simulation in which only one cyclotron harmonic was excited.

We then analysed a novel spectrally structured ICE feature in the range 500 MHz to 900 MHz, which also exhibited chirping on sub-microsecond timescales. The spectral peaks of this feature correspond to harmonics l of the proton cyclotron frequency ω_{cp} at the outer midplane edge, where $l = 20$ to 36 . This frequency range exceeded estimates of ω_{LH} from our previous analysis. This new feature was time-shifted with respect to the brighter lower-frequency chirping ICE feature previously analysed. By carrying out bispectral analysis of the measured KSTAR fields and of the field amplitudes output from the PIC simulations, we showed that the new, fainter, higher-frequency chirping ICE feature, was driven by nonlinear wave coupling between different neighbouring spectral peaks in the lower-frequency ICE feature. This reinforces the identification of the MCI as the plasma physics process underlying proton harmonic ICE from KSTAR, while providing a novel instance of nonlinear wave coupling on very fast timescales.

We then turned our attention to the “steady state” deuterium ICE which is observed in multiple KSTAR pulses, usually $\sim 50\mu s - 100\mu s$ prior to the ELM crash and in some KSTAR pulses, immediately before the chirping features described above. Using the linear analytical theory of the MCI, energetic particle orbit studies, and first principles PIC simulations, we were able to provide an explanation for the origin of steady-state ICE at multiple deuterium cyclotron harmonics. The collective relaxation of a minority energetic population of NBI deuterons in two PIC simulations generated electric and magnetic field oscillations whose power spectra substantially resemble the measured ICE spectra. Some low harmonic peaks in one

simulation frequency spectrum were not detected in its experimental counterpart. A probable explanation for this is that the S_{11} return loss of the Bowtie antenna used to measure this RF signal was close to 0dB, implying very high reflectivity in this low frequency range.

We have now explained the main features of two distinct types of ICE from KSTAR plasmas: steady-state ICE due to NBI deuterons; and highly dynamic chirping ICE due to fusion-born protons. It seems appropriate to conjecture that these two types of ICE events naturally follow each other during ELM crashes, at least in KSTAR. The “steady state” NBI ICE exists before the ELM crash and can be reproduced with PIC simulations using a background electron number density similar to that found at the top of the pedestal. As the pedestal collapses through multiple filament bursts, the local density changes, so the ICE is no longer “steady state” and instead chirps down, which can be reproduced by running PIC simulations with a range of electron number densities. We investigated this conjecture by running two additional simulations with different minority ion concentrations, initialising both minority NBI deuterons and minority fusion-born protons using ring-beam velocity distributions, the latter being time delayed with respect to the former. We found that in both simulations, the fusion-born proton ICE was much more spectrally intense than the NBI deuteron ICE, offering a potential explanation as to why steady-state ICE in steps of the deuteron cyclotron frequency, and highly dynamic downward frequency chirping ICE in steps of the proton cyclotron frequency, are not observed at the same time during the KSTAR ELM crash process. This tentative result does not explain why the ICE spacing is ω_{cD} before the crash, and sometimes changes to ω_{cp} during the crash. For the ICE spacing to change, the driving populations of energetic particles must be different - NBI deuterons prior to the crash, and fusion born protons during the crash. We should note that downward frequency chirping with spacing equal to ω_{cD} is also observed during some KSTAR plasmas. Further multiple ion species simulations may shed light on the existence of the side-band features observed in some KSTAR plasmas that exhibit downward frequency chirping in steps of the proton cyclotron frequency. Reconciling these three ICE observations and putting them into the context of the overall ELM crash cycle is the subject of future work.

8.2 JET and AUG ICE

This chapter started by first examining the MCI instability in relation to the ICRF ICE JET observed in multiple JET plasma pulses reported in Refs. [Jacquet et al., 2011; McClements et al., 2018]. Simulations of the MCI in a proton plasma with a minority helium-3 ring beam population were run, and successfully reproduced the experimentally observed ICE feature at the ^3He fundamental. We then went on to examine how different wave propagation angles θ and the inclusion of a finite parallel drift $v_{\parallel 0}$ in the ^3He ring beam distribution affect the simulation dynamics. It was found that the simulations were somewhat sensitive to a small change in θ from 90° to 89° , but that the inclusion of a finite parallel drift $v_{\parallel 0}$ had little effect on the dynamics of simulations with $\theta = 89^\circ$. We found that for these plasma parameters, increasing θ to 92° , led to a predominantly electrostatic instability, and, increasing it further beyond 92° , meant the instability was no longer excited. Future work could investigate the role of a finite $v_{\parallel 0}$ in simulations with θ further away from 90° , as the effects of parallel dynamics will become more important, provided we are still close enough to $\theta = 90^\circ$ for the MCI to be excited strongly.

We then analysed the results of a single MCI simulation of the core FP ICE observed at the fundamental proton cyclotron frequency in AUG deuterium plasmas [Ochoukov et al., 2018]. We approximated the expected proton spherical shell distribution function as a proton ring beam, on the basis that only a narrow region of phase space contributes to the excitation of ICE via the MCI. Using bicoherence analysis, it was shown the linear stable modes with $l > 5$, on both the forward and backward propagating branches of the dispersion relation, couple together nonlinearly to produce many additional spectral features, including the experimentally observed $l = 1$ mode and other linearly stable modes with $l \leq 5$. This $l = 1$ feature was observed to grow in amplitude as the simulation progresses deeper into the nonlinear re-energisation regime, and the strength of the nonlinear coupling was observed to follow a similar pattern.

In both the JET and AUG relevant simulations described above, the nonlinear aspects of the MCI and ICE excitation were shown to be indispensable to the interpretation of the observed ICE phenomena in both devices. Future devices such as ITER would benefit from an effort to detect high cyclotron harmonics, such as those observed in KSTAR (see Refs. [Thatipamula et al., 2016; Kim et al., 2018] and chapter 4). Provided the temporal resolution is sufficiently high, it would then be possible to use bicoherence analysis as a first step to quantifying the nonlinear interactions which may lead to low l modes, which could then be used to confirm

the hypothesis set out in chapter 5.

8.3 Core ICE

In this chapter, we performed the first PIC simulations of the MCI due to the collective relaxation of minority energetic protons modelled using an isotropic spherical shell distribution. We analysed six simulations of the MCI under these conditions, varying the velocity spread in each one, as well as six additional simulations of the MCI in which the minority protons were initialised using ring-beam distribution functions of varying perpendicular velocity spread. The MCI was excited in all cases, and the spherical shell simulations took around twice as long to saturate as their ring beam counterparts, with around ten times less energy transfer from the minority ions to the bulk plasma and electromagnetic fields.

The mode structure in both sets of simulations was found to be qualitatively similar, and, in each case, increasing the velocity spread led to the most spectrally intense mode gradually shifting from the fourth proton cyclotron harmonic, to the fifth, and finally to the sixth. This suggests that by observing the mode structure in experiments, one may be able to deduce the spread of the minority ion distribution, an important parameter as it determines how much energy is transferred to the bulk plasma and electromagnetic fields due to the MCI. The nonlinear aspects of each simulation were discussed and many significant nonlinear wave-wave couplings were identified. In general, there are many more wave-wave interactions in the ring-beam simulations than in the spherical shell simulations, however, the strongest nonlinearly driven modes of practical interest, i.e. those that contribute significantly to the ICE signal, are present among both sets of simulations. In particular, both sets of simulations exhibit a strong nonlinearly driven $(k, l) \approx (8.7, 8)$ mode which does not lie along the magnetoacoustic dispersion branch, and in the case of the ring-beam simulation with zero velocity spread, contributes to approximately half of the total intensity of the eighth proton cyclotron harmonic. Other nonlinear couplings gave rise to modes above the lower hybrid frequency ω_{LH} , and in the ring-beam simulations, some of these modes had intensities comparable to low intensity linearly unstable modes. This demonstrates how indispensable the nonlinear physics is when simulating ICE and interpreting experimental observations. The key to identifying mode couplings was to first fulfil the wavenumber matching criterion, followed by the frequency matching criterion. These two requirements highlight that a modest experimental effort to detect both the perpendicular wavenumber

and high frequency ion cyclotron harmonics would enable us to better understand the measured frequency spectrum, and hence the character of the energetic ion distribution function.

The similarity between: the variation of energy density with spread, the linearly excited mode structure, and the nonlinear characteristics of both sets of simulations, suggests that a ring-beam velocity distribution for the minority energetic ions serves as a close approximation to an isotropic spherical shell distribution, provided the velocity spread is not too large. From a resource perspective, this is crucial. The diagnostics with which to measure core ICE in tokamaks are becoming more widespread, and the computing resources with which to simulate it are becoming increasingly more sophisticated; we are thus at a juncture in which it will soon be feasible for PIC simulations of the MCI to be used for predictive modelling of tokamak plasma phenomena, as opposed to only interpretive modelling. The “cheaper” ring-beam simulations offer a way to realise this.

8.4 Preliminary simulations of helium ash pumping in JET core plasmas

In this short section, we analysed preliminary results of PIC simulations pertaining to helium ash pumping in JET core plasmas. We identified a novel mechanism in which perpendicular deuteron NBI can be used to selectively target helium ash in the core, transferring a considerable amount of perpendicular momentum to a subset of the ash population, which may alter its orbit in a way which carries them out of the core plasma. Only low temperature helium ash populations were affected by this process, meaning that it may be possible to modify existing NBI systems, or implement new ones, to pump out this core ash, with negligible consequences for the helium ions that are still energetic and can be used for plasma heating. The simulations presented in this section are the subject of ongoing research, and further work is underway to quantify the effect of NBI beams on helium ash populations in core tokamak plasmas.

Appendix A

Hybrid version of EPOCH

A.1 Introduction to hybrid codes

I have begun work on a hybrid counterpart to the 1D3V version of the EPOCH PIC code detailed in chapter 4. This code can be used in conjunction with the PIC version of EPOCH, requiring only a compiler flag to be switched on, along with the specification of some additional input parameters which are discussed in Sec. A.3.1. In this context, a “hybrid” code is one in which the electrons are treated as a massless neutralising fluid, while the ions are represented as computational macro-particles like in a regular PIC code. This approximation is relevant to phenomena in which the typical length scales are larger than the ion inertial length, and the time scales are of the order of the ion gyro-period [Winske et al., 2003]. Hybrid codes allow us to study plasma phenomena occurring on long time-scales, whilst retaining the full gyro-motion of the ions. Of course, one must take care to avoid neglecting any electron effects that play a role in the structure of waves supported predominantly by the thermal motion of the ions. An example of this is given in Sec. 4.2, where it was observed that the lower hybrid frequency, which depends on the electron mass (see Eq. 1.15), curtails the number of modes that are available for excitation via the energetic-ion driven magnetoacoustic cyclotron instability (MCI). Hybrid PIC codes have been successfully applied to a wide range of plasma physics phenomena, particularly those relevant to space and solar wind physics, see Ref. [Winske et al., 2003] and references therein. Recently, hybrid codes have been successfully applied to MCF physics, such as the formation of filamentary structures “plasma blobs” in the edge region of tokamak plasmas [Gingell, 2013; Gingell et al., 2012, 2014, 2013], and simulations of Ion cyclotron emission (ICE) via the MCI in a range of tokamak operating regimes [Carbajal et al., 2014; Carbajal, 2015; Carbajal et al.,

2017; Dendy et al., 2017; Reman et al., 2016]. In Sec. A.2, we discuss the equations used to advance the electromagnetic fields in Hybrid EPOCH, and follow this by a discussion of the numerical scheme in Sec. A.3. Most of the technical aspects of Hybrid EPOCH, such as the particle shape functions, the particle loading, and the I/O, are identical to that of the standard PIC version of EPOCH. As such, only major differences are discussed in this appendix, and we refer the reader to chapter 2 for more details.

It should be stressed to the reader that at the time of writing, the 1D3V version of Hybrid EPOCH is not yet functional, and suffers from poor energy conservation. This appendix is here to serve as a record of the code development thus far, to aid in the future development of Hybrid EPOCH.

A.2 Hybrid field equations

We start with the electron momentum equation, and let the $m_e \frac{dv_e}{dt}$ term on the left side of the equation equal 0

$$0 = en_e \left(\mathbf{E} + \frac{\mathbf{V}_e \times \mathbf{B}}{c} \right) - \nabla \cdot \mathbf{P}_e. \quad (\text{A.1})$$

Here, \mathbf{E} is the electric field, \mathbf{B} is the magnetic field, \mathbf{V}_e is the electron fluid velocity, and \mathbf{P}_e is the electron pressure tensor. We assume quasi-neutrality such that $n_e = \sum_j^N Z_j n_j$, where n_e is the electron density, N is the total number of ion species, n_j is the number density of an ion belonging to species j , and Z_j is the atomic number of an ion belonging to species j . To simplify the algebra, we will assume that there is only one ion species in the following derivation, it is then straightforward to extend the equations to multiple ion species. Here we will take the pressure as a scalar such that $\mathbf{P}_e = p_e \mathbf{I}$ where \mathbf{I} is the unit dyadic and p_e is the scalar electron pressure. We set $p_e = n_e k_b T_e$, where T_e is assumed constant throughout the simulation.

In its present form, the hybrid version of EPOCH assumes no resistivity. This can easily be implemented later by adding the term $en_e \eta \cdot \mathbf{J}$ to the right side of Eq. A.1, where η is the resistivity tensor and \mathbf{J} is the total current; along with a slight modification to the particle push routine to add $-e\eta \cdot \mathbf{J}$ to acceleration term in ion equations.

We use Ampère’s law in low frequency limit

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} = \mu_0 q_i n_i (\mathbf{V}_i - \mathbf{V}_e), \quad (\text{A.2})$$

where \mathbf{V}_i is the ion bulk velocity

$$\mathbf{V}_i = \frac{\sum_{j=1}^N (Z_j n_j \mathbf{u}_j)}{\sum_{j=1}^N (Z_j n_j)}. \quad (\text{A.3})$$

N is the number of ion species, \mathbf{u}_j and n_j are the bulk velocity and number density of species j . The bulk velocity of an ion belonging to species j is

$$\mathbf{u}_j = \frac{\sum_{k=1}^M \Delta w_k(\mathbf{r}) \mathbf{v}_k}{\sum_{k=1}^M \Delta w_k(\mathbf{r})} \quad (\text{A.4})$$

Where $\Delta w_k(\mathbf{r})$ is the contribution of the ion to the grid cell at position \mathbf{r} . Combining to eliminate variables gives

$$\mathbf{E} = \frac{1}{\mu_0 e n_e} ((\nabla \times \mathbf{B}) \times \mathbf{B}) - \mathbf{V}_i \times \mathbf{B} - \frac{k_b T_e}{e n_e} \nabla n_e, \quad (\text{A.5})$$

where n_e is given by the quasi-neutrality condition. Equation A.5 and Faraday’s law are the equations used to evolve the electromagnetic fields in Hybrid EPOCH. The position and velocity are updated in the same way as the PIC version of EPOCH, which is detailed in chapter 2. To summarise, the only change to the equations is that Ampère’s law has been swapped out for the the generalised Ohm’s law derived here. Note that Eq. A.5 is independent of time, only grid quantities defined at the appropriate times are required to evolve the electric field.

A.3 Numerical implementation

Two variations of a similar algorithm have been implemented in Hybrid EPOCH, and their usage is controlled by a compiler flag. In both variants of the algorithm, the particle’s positions and velocities are staggered in time, such that the particles positions are defined at time level “ N ”, and the particle’s velocities are defined at time level “ $N - \frac{1}{2}$ ”. The particle push is identical to that used in the PIC version of EPOCH, the only difference is that we no longer need the extra half timestep position update to get the current using the Villasenour-Buneman method [Villasenor and Buneman, 1992], and there are two additional function calls to get n_i and \mathbf{V}_i on the grid. In the following, the functions: **F**, **G**, **H**, and **L** denote components of: the Lorentz force law used to update the velocities, the equation of motion, Faraday’s

law, and the generalised Ohm's law. The default algorithm can be summarised by the following steps:

1. Update the particle velocity by a full timestep: $\mathbf{v}^{N+1/2} = \mathbf{v}^{N-1/2} + \mathbf{F}(\mathbf{E}^N, \mathbf{B}^N)$.
2. Update the particle position by a full timestep: $\mathbf{x}^{N+1} = \mathbf{x}^N + \mathbf{G}(\mathbf{v}^{N+1/2})$.
 → Gather the ion bulk velocity \mathbf{V}_i at time level $N + 1/2$, as well as the total ion density n_i at time levels $N + 1/2$ and $N + 1$.
3. Update the magnetic field by half a timestep: $\mathbf{B}^{N+1/2} = \mathbf{B}^N + \mathbf{H}(\mathbf{E}^N)$.
4. Update the electric field by half a timestep: $\mathbf{E}^{N+1/2} = \mathbf{L}(\mathbf{B}^{N+1/2}, n_i^{N+1/2}, \mathbf{V}_i^{N+1/2})$.
5. Update the magnetic field by a further half timestep: $\mathbf{B}^{N+1} = \mathbf{B}^{N+1/2} + \mathbf{H}(\mathbf{E}^{N+1/2})$.
6. Use 4th-order Bashford-Adams extrapolation to obtain the ion bulk velocity at time level $N + 1$ (details below): \mathbf{V}_i^{N+1} .
7. Update the electric field by a further half timestep: $\mathbf{E}^{N+1} = \mathbf{L}(\mathbf{B}^{N+1}, n_i^{N+1}, \mathbf{V}_i^{N+1})$.

The bulk velocity at time level $N+1$ is unknown, once the simulation has advanced at least two timesteps, this is estimated using 4th-order Bashford-Adams extrapolation

$$\mathbf{V}_i^{N+1} = 2\mathbf{V}_i^{N+1/2} - \frac{3}{2}\mathbf{V}_i^{N-1/2} + \frac{1}{2}\mathbf{V}_i^{N-3/2}. \quad (\text{A.6})$$

For the first timestep, we necessarily use 2nd-order extrapolation

$$\mathbf{V}_i^{N+1} = \frac{3}{2}\mathbf{V}_i^{N+1/2} - \frac{1}{2}\mathbf{V}_i^{N-1/2}. \quad (\text{A.7})$$

The second variant of this algorithm uses 4th-order Runge-Kutta subcycling, in which the magnetic field is advanced from time level N to time level $N + 1$ by dividing the update into θ smaller timesteps of duration $\Delta t' = \Delta t/\theta$. This has been shown to ameliorate the effects of the short wavelength whistlers that plague most hybrid codes [Terasawa et al., 1986; Winske et al., 2003]. We will denote a subcycled time-level by $s = N/\theta$. The magnetic field at time level $N + s$ is given by

$$\mathbf{B}^{N+s} = \mathbf{B}^N + \frac{\Delta t'}{6} (\mathbf{K}_1^N + 2\mathbf{K}_2^N + 2\mathbf{K}_3^N + \mathbf{K}_4^N), \quad (\text{A.8})$$

where

$$\mathbf{K}_1^N = -\nabla \times \mathbf{L}(\mathbf{B}^N),$$

$$\begin{aligned}
\mathbf{K}_2^N &= -\nabla \times \mathbf{L} \left(\mathbf{B}^N + \frac{\Delta t'}{2} \mathbf{K}_1^N \right), \\
\mathbf{K}_3^N &= -\nabla \times \mathbf{L} \left(\mathbf{B}^N + \frac{\Delta t'}{2} \mathbf{K}_2^N \right), \\
\mathbf{K}_4^N &= -\nabla \times \mathbf{L} \left(\mathbf{B}^N + \Delta t' \mathbf{K}_3^N \right),
\end{aligned} \tag{A.9}$$

and \mathbf{L} is the function we defined earlier to represent the electric field update, and takes $n_i^{N+1/2}$ and $\mathbf{V}_i^{N+1/2}$ as inputs throughout the entire magnetic field update. This alternative algorithm can be switched on by setting the ‘‘SUBCYCLING’’ compiler flag in the Makefile (in addition to the ‘‘HYBRID’’ flag), and works in the following way:

1. Update the particle velocity by a full timestep: $\mathbf{v}^{N+1/2} = \mathbf{v}^{N-1/2} + \mathbf{F}(\mathbf{E}^N, \mathbf{B}^N)$.
2. Update the particle position by a full timestep: $\mathbf{x}^{N+1} = \mathbf{x}^N + \mathbf{G}(\mathbf{v}^{N+1/2})$.
→ Gather the ion bulk velocity \mathbf{V}_i at time level $N + 1/2$, as well as the total ion density n_i at time levels $N + 1/2$ and $N + 1$.
3. Update the magnetic field by a full timestep using 4th order Runge-Kutta subcycling: $\mathbf{B}^{N+1} = \mathbf{B}^N +$ Loop over θ iterations:
→ $\mathbf{B}^{N+s} = \mathbf{B}^{N+s} = \mathbf{B}^N + \frac{\Delta t'}{6} (\mathbf{K}_1^N + 2\mathbf{K}_2^N + 2\mathbf{K}_3^N + \mathbf{K}_4^N)$ where $\mathbf{K}_{\{1..4\}}$ are given by Eq. A.9.
4. Use 4th order Bashford-Adams extrapolation to obtain the ion bulk velocity at time level $N + 1$: \mathbf{V}_i^{N+1} .
5. Update the electric field by a full timestep: $\mathbf{E}^{N+1} = \mathbf{L}(\mathbf{B}^{N+1}, n_i^{N+1}, \mathbf{V}_i^{N+1})$.

Note that both schemes require only one loop through the particles. Regardless of which scheme is used, both schemes still suffer from a lack of energy conservation due to short wavelength whistler modes. This can be ameliorated by smoothing the grid variables \mathbf{B} , \mathbf{E} , \mathbf{V}_i , and n_i at each iteration. For a given quantity H located at grid point p a weighted average is calculated

$$\overline{H}_p = \frac{1}{4}H_{p-1} + \frac{1}{2}H_p + \frac{1}{4}H_{p+1}. \tag{A.10}$$

The final value of H is given by

$$H_p = (1 - \alpha_s) H_p + \alpha_s \overline{H}_p, \quad (\text{A.11})$$

where α_s is a user specified smoothing parameter controlling the level of numerical diffusion.

A.3.1 Additions to the EPOCH input deck

A new block called “hybrid” has been added to the EPOCH input deck. There are only three new values for the user to set:

Te_x - The electron temperature in the x-direction, the only relevant direction in the current 1D3V version of EPOCH. The default value is 0, which is equivalent to not including the pressure gradient term in Eq. A.1.

alpha - The smoothing parameter defined above. The default value is 0, but, if this is not set, the program provides a warning telling the user that this is ill advised.

rk4_steps - The number of subcycling steps to do if the code was compiled using the “SUBCYCLING” flag. The default and value is 1, which is also the minimum value allowed. If the code was not compiled with the “SUBCYCLING” option, this entry is ignored.

All default values of these three quantities are all valid under the right physical conditions, meaning Hybrid EPOCH can in theory be run with no changes to the input deck, and only one change to the Makefile. Of course, this may have implications for the physics, if electron pressure is important, or if the simulation is sensitive to short wavelength whistler modes.

A.4 Summary and future

In this appendix we have provided documentation pertaining to the implementation of a hybrid 1D3V version of the EPOCH PIC code, in which the electrons are treated as a massless neutralising fluid and the ions are treated as particles. This version of the code is finished, but still suffers from poor energy conservation and instability. This will be addressed in the future. Once the 1D3V version of the code is numerically stable, it is straightforward to extend the code to 2D3V and 3D3V, and to add other physics effects such as finite plasma resistivity and finite electron mass.

Appendix B

Mapping between experimental and simulation harmonics

The following shows an example procedure for computing the empirical relationship between the electron particle density n_e and time t which has been calculated for three KSTAR pulses shown in Figs. 4.5 and 4.7. This example corresponds to KSTAR pulse 11462, and the result of this analysis is shown in the left panel of 4.5. The Fourier power is plotted using a \log_{10} scale in all figures. The general idea is to compare experimentally observed harmonics, with those observed in the simulation, noting the value of n_e in the simulations which excite waves at the frequencies observed experimentally. It should be noted that there is significantly more resolution in frequency space experimentally, than there is computationally. The procedure is as follows.

Step 1 A corresponding experimental and computational frequency range are identified and sectioned off from the rest of the plot for comparison. This is shown in Fig. B.1.




Figure B.1: Top: Experimental spectrogram showing the power in the 17th proton cyclotron harmonic as a function of frequency and time. Bottom: The result of several independent simulations of the MCI which shows the power in the 17th harmonic as a function of frequency and n_e .

Step 2 Only the most powerful simulation modes are considered in the construction of the $n_e(t)$ points for a given harmonic. For instance, the section of of the bottom plot in Fig. B.1 with $1.3 \times 10^{19} \text{m}^3 \leq n_e \leq 1.4 \times 10^{19} \text{m}^3$ is not considered to contribute to the Fourier power in its experimental counterpart, as the power in the simulation is comparatively weaker than the power in the simulation at $n_e > 1.4 \times 10^{19} \text{m}^3$. The experimental spectrogram is also curtailed along the time axis, so only the most powerful sections of the mode remain. This is shown in Fig. B.2.

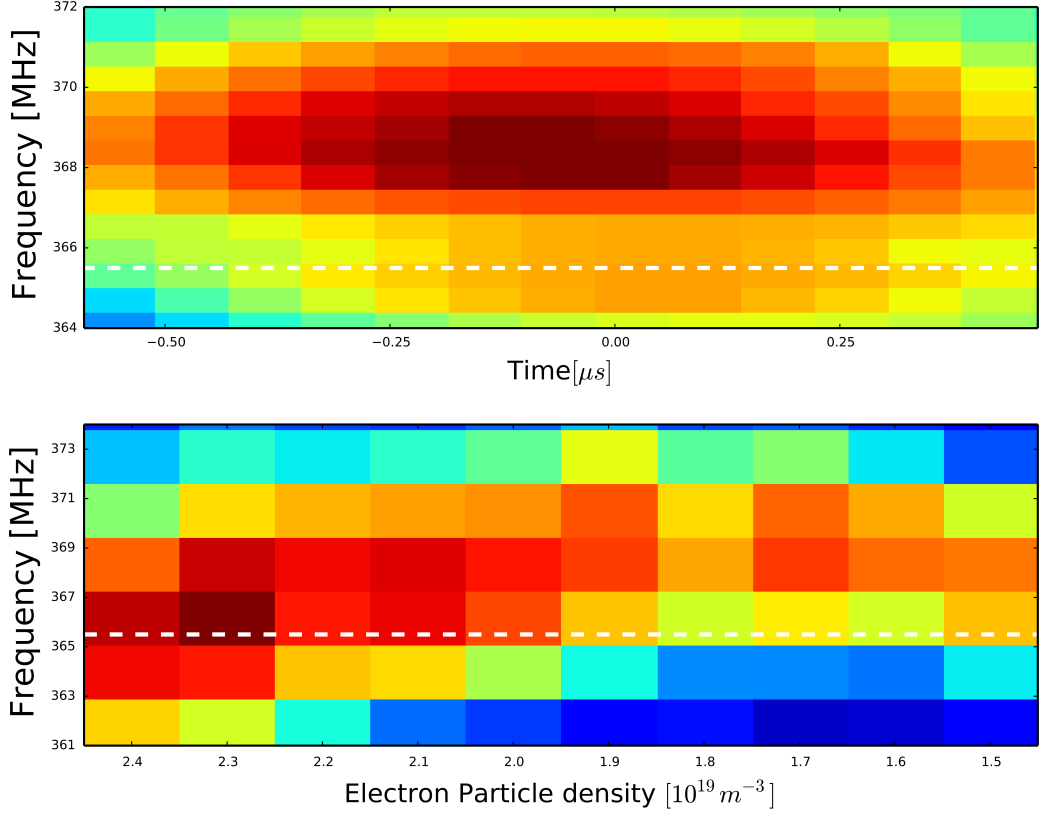


Figure B.2: Top: Experimental spectrogram showing the power in the 17th harmonic as a function of a narrow range of frequency and time. Bottom: The result of several independent simulations of the MCI which shows the power in the 17th harmonic as a function of a narrow range of frequency and n_e .

Step 3 The feature is then divided into N data points of equal width in n_e and t .

The lower panel in Fig. B.2 shows 10 different values of n_e corresponding to the experimental feature. Thus, the time axis of the upper plot is divided into 10 equidistant time points. In Fig. B.2, the start of the experimental feature is at $t_i \simeq -0.59 \mu s$, this corresponds to $n_e = 2.4 \times 10^{19} m^3$ in the simulation. The end of the feature is at $t_f \simeq 0.47 \mu s$, this corresponds to $n_e = 1.5 \times 10^{19} m^3$ in the simulation. Dividing the experimental feature into 10 time points, gives a time width $\Delta t = \frac{t_f - t_i}{N-1} \simeq \frac{0.47 - (-0.59)}{9} \mu s \simeq 0.117 \mu s$. Thus the point $n_e = 2.3 \times 10^{19} m^3$ corresponds to a time $t_i + \Delta t \simeq (-0.59 + 0.117) \mu s \simeq -0.47 \mu s$, the point $n_e = 2.2 \times 10^{19} m^3$ corresponds to a time $t_i + 2\Delta t \simeq (-0.59 + (2 \times 0.117)) \mu s \simeq -0.36 \mu s$, and so on. The error in time for each data point is taken to be Δt , always rounded up to 2 decimal places, while the error in density is taken to be the difference between 2 successive density

points, which remains constant at $0.1 \times 10^{19} \text{m}^3$. Note, while Δt is calculated to 3 decimal places, all resulting time points are rounded up to 2 decimal places. Fig. B.3 demonstrates the above pictorially.




Figure B.3: Top: Experimental spectrogram showing the power around the 17th harmonic as a function of a narrow range of frequency and time. The black boundary to the left of the letter indicates the associated time point, the black boundaries to the left of t_i and t_f indicate the first and last time points respectively. Bottom: Power in the simulations around the 17th harmonic as a function of a narrow range of frequency and n_e . Δt indicates the width between successive time points, the letters (a) to (j) denote a mapping between n_e and t .

Step 4 As n_e decreases, and the lower hybrid frequency ω_{LH} starts to play a more dominant role in the simulation dynamics, the number of simulations, and hence the number of n_e points corresponding to a given experimental feature drastically reduces. A similar procedure to the one described in step 3 is followed, leading to larger error bars in time. An example of this is shown in Fig. B.4, in which there are only two values of n_e corresponding to a given feature in the experimental spectrogram. This experimental feature is around the same width in time as the others, (due to the finite window used in the Fourier transform), as such, the time error bars in this region of (t, n_e) parameter space are much larger. In this case, the time points t_i and t_f are taken to be the start and end of the very brightest part of the feature. The error in t_i , which we call Δt_i , is estimated as either the difference in time between t_i and the left edge of the feature, or the spacing between t_i and t_f , whichever is larger. The error in t_f , which we call Δt_f , is found in the same way, but using the right edge of the feature. In this example, Δt_i is estimated as the difference in time between t_i and the left edge of the feature, while Δt_f is estimated as the spacing between t_i and t_f , as these provide the largest error estimate in both cases. These points are marked on Fig. B.4.

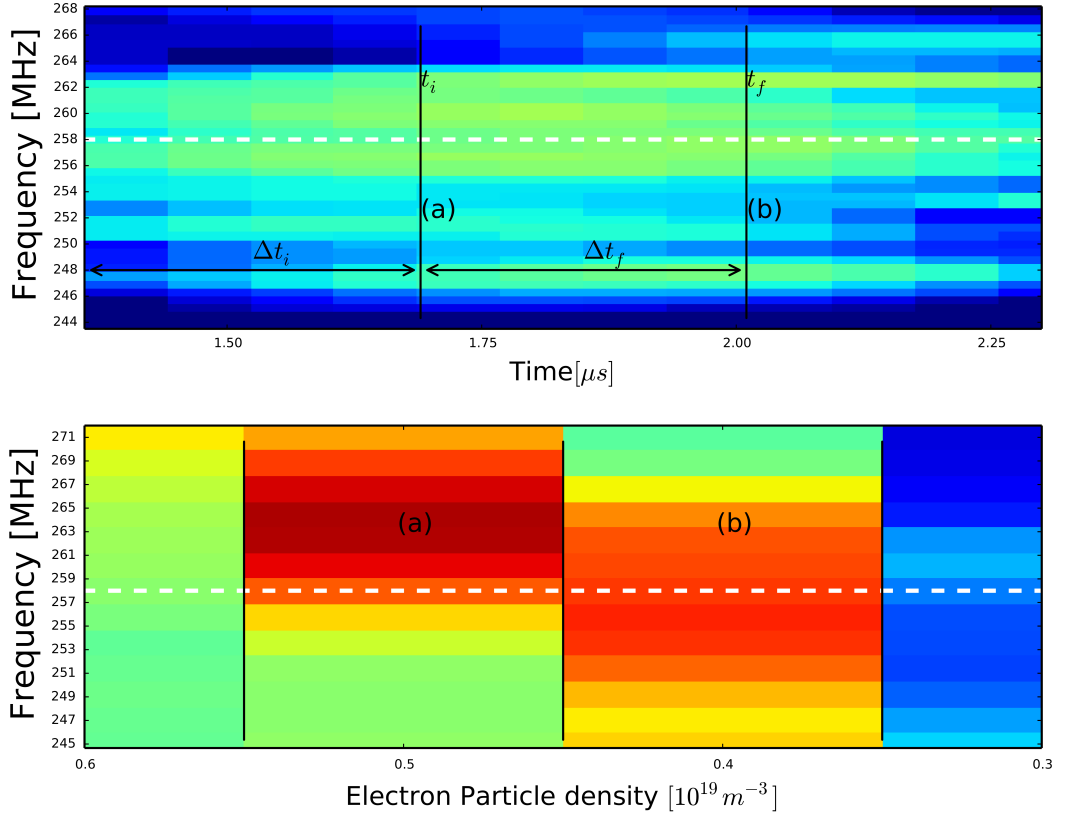


Figure B.4: Top: Experimental spectrogram showing the power in the 12th harmonic as a function of frequency and time. Bottom: Power around the corresponding 12th harmonic as a function of frequency and n_e . t_i and t_f denote the time points corresponding to $n_e = 0.5 \times 10^{19} \text{m}^{-3}$ and $n_e = 0.4 \times 10^{19} \text{m}^{-3}$ respectively. Δt_i and Δt_f indicate the estimated errors in t_i and t_f respectively. The labels (a) and (b) denote a mapping between n_e and t .

Step 5 As n_e decreases further, a value is reached at which only one mode in the simulation can be excited. In this case, the value of t for a given n_e is estimated to be the centre point of the bright region of the corresponding experimental feature. The error in this is estimated as half the width of the feature, often yielding the largest error in time from the entire mapping procedure. Fig. B.5 demonstrates this. As this corresponds to just one value of n_e , there is only one panel showing the relevant section of the experimental spectrogram.




Figure B.5: Experimental spectrogram showing the power in the 11th harmonic as a function of frequency and time. t denotes the time point which in this case corresponds to $n_e = 0.5 \times 10^{19} \text{m}^{-3}$. The approximate width of the bright region of the harmonic is denoted by $2\Delta t$, where Δt is the estimated error.

Bibliography

- T D Arber, K Bennett, C S Brady, A Lawrence-Douglas, M G Ramsay, N J Sircombe, P Gillies, R G Evans, H Schmitz, A R Bell, and C P Ridgers. Contemporary particle-in-cell approach to laser-plasma modelling. *Plasma Physics and Controlled Fusion*, 57(11):1–26, nov 2015.
- L.G. Askinazi, A.A. Belokurov, D.B. Gin, V.A. Kornev, S.V. Lebedev, A.E. Shevelev, A.S. Tukachinsky, and N.A. Zhubr. Ion cyclotron emission in NBI-heated plasmas in the TUMAN-3M tokamak. *Nuclear Fusion*, 58(8):082003, 2018. URL <http://stacks.iop.org/0029-5515/58/i=8/a=082003>.
- V. S. Belikov and Ya. I. Kolesnichenko. Magnetoacoustic cyclotron instability in a thermonuclear plasma. *Sov. Phys. Tech. Phys.*, 20:1146, 1976.
- T. S. Bird. Definition and misuse of return loss [report of the transactions editor-in-chief]. *IEEE Antennas and Propagation Magazine*, 51(2):166–167, April 2009. ISSN 1045-9243. doi: 10.1109/MAP.2009.5162049.
- Charles K. Birdsall and A. Bruce Langdon. *Plasma Physics Via Computer*. McGraw-Hill, Inc., New York, NY, USA, 1985. ISBN 0070053715.
- O Buneman. *Computer Space Plasma Physics: Simulations Techniques and Software*. Tokyo: Terra Scientific, 1993.
- R. A. Cairns. *Plasma Physics*, pages 62–92. Dordrecht: Kluwer Academic Publishers Group, 1985.
- R. A. Cairns. *Plasma Physics an Introductory Course (edited by Richard Dendy)*, pages 397–399. Cambridge university press - Cambridge, 1996.
- L. Carbajal, R. O. Dendy, S. C. Chapman, and J. W. S. Cook. Linear and nonlinear physics of the magnetoacoustic cyclotron instability of fusion-born ions in relation to ion cyclotron emission. *Physics of Plasmas*, 21(1):012106, 2014. doi: 10.1063/1.4861866. URL <https://doi.org/10.1063/1.4861866>.

- L. Carbajal, R. O. Dendy, S. C. Chapman, and J. W. S. Cook. Quantifying Fusion Born Ion Populations in Magnetically Confined Plasmas using Ion Cyclotron Emission. *Phys. Rev. Lett.*, 118:105001, Mar 2017. doi: 10.1103/PhysRevLett.118.105001. URL <https://link.aps.org/doi/10.1103/PhysRevLett.118.105001>.
- Leopoldo Carbajal. *Transport in turbulent plasmas at the interface between different levels of description*. PhD thesis, The University of Warwick, 2015.
- S. Cauffman, R. Majeski, K.G. McClements, and R.O. Dendy. Alfvénic behaviour of alpha particle driven ion cyclotron emission in TFTR. *Nuclear Fusion*, 35(12):1597, 1995. URL <http://stacks.iop.org/0029-5515/35/i=12/a=I22>.
- B. Chapman, R.O. Dendy, K.G. McClements, S.C. Chapman, G.S. Yun, S.G. Thatipamula, and M.H. Kim. Sub-microsecond temporal evolution of edge density during edge localized modes in KSTAR tokamak plasmas inferred from ion cyclotron emission. *Nuclear Fusion*, 57(12):124004, 2017. URL <http://stacks.iop.org/0029-5515/57/i=12/a=124004>.
- B. Chapman, R.O. Dendy, S.C. Chapman, K.G. McClements, G.S. Yun, S.G. Thatipamula, and M.H. Kim. Nonlinear wave interactions generate high-harmonic cyclotron emission from fusion-born protons during a kstar elm crash. *Nuclear Fusion*, 58(9):096027, 2018. URL <http://stacks.iop.org/0029-5515/58/i=9/a=096027>.
- I T Chapman. Controlling sawtooth oscillations in tokamak plasmas. *Plasma Physics and Controlled Fusion*, 53(1):013001, nov 2010. doi: 10.1088/0741-3335/53/1/013001. URL <https://doi.org/10.1088%2F0741-3335%2F53%2F1%2F013001>.
- I.T. Chapman, J. Adamek, R.J. Akers, S. Allan, L. Appel, O. Asunta, M. Barnes, N. Ben Ayed, T. Bigelow, W. Boeglin, J. Bradley, J. Brnner, P. Cahyna, M. Carr, J. Caughman, M. Cecconello, C. Challis, S. Chapman, J. Chorley, G. Colyer, N. Conway, W.A. Cooper, M. Cox, N. Crocker, B. Crowley, G. Cunningham, A. Danilov, D. Darrow, R. Dendy, A. Diallo, D. Dickinson, S. Diem, W. Dorland, B. Dudson, D. Dunai, L. Easy, S. Elmore, A. Field, G. Fishpool, M. Fox, E. Fredrickson, S. Freethy, L. Garzotti, Y.C. Ghim, K. Gibson, J. Graves, C. Gurl, W. Guttenfelder, C. Ham, J. Harrison, D. Harting, E. Havlickova, J. Hawke, N. Hawkes, T. Hender, S. Henderson, E. Highcock, J. Hillesheim, B. Hnat, J. Holgate, J. Horacek, J. Howard, B. Huang, K. Imada, O. Jones, S. Kaye, D. Keeling, A. Kirk, I. Klimek, M. Kocan, H. Leggate, M. Lilley, B. Lipschultz, S. Lisgo, Y.Q.

Liu, B. Lloyd, B. Lomanowski, I. Lupelli, G. Maddison, J. Mailloux, R. Martin, G. McArdle, K. McClements, B. McMillan, A. Meakins, H. Meyer, C. Michael, F. Militello, J. Milnes, A.W. Morris, G. Motojima, D. Muir, E. Nardon, V. Naulin, G. Naylor, A. Nielsen, M. O'Brien, T. O'Gorman, Y. Ono, H. Oliver, S. Pamela, L. Pangione, F. Parra, A. Patel, W. Peebles, M. Peng, R. Perez, S. Pinches, L. Piron, M. Podesta, M. Price, M. Reinke, Y. Ren, C. Roach, J. Robinson, M. Romanelli, V. Rozhansky, S. Saarelma, S. Sangaroon, A. Saveliev, R. Scannell, A. Schekochihin, S. Sharapov, R. Sharples, V. Shevchenko, S. Silburn, J. Simpson, J. Storrs, Y. Takase, H. Tanabe, H. Tanaka, D. Taylor, G. Taylor, D. Thomas, N. Thomas-Davies, A. Thornton, M. Turnyanskiy, M. Valovic, R. Vann, N. Walkden, H. Wilson, L.V. Wyk, T. Yamada, S. Zoletnik, MAST, and MAST Upgrade Teams. Overview of MAST results. *Nuclear Fusion*, 55(10):104008, 2015. URL <http://stacks.iop.org/0029-5515/55/i=10/a=104008>.

Francis F. Chen. *Plasma physics and controlled fusion*, pages 30–33. Springer, 1984.

J. W. S. Cook, S. C. Chapman, and R. O. Dendy. Electron Current Drive by Fusion-Product-Excited Lower Hybrid Drift Instability. *Phys. Rev. Lett.*, 105:255003, Dec 2010. doi: 10.1103/PhysRevLett.105.255003. URL <https://link.aps.org/doi/10.1103/PhysRevLett.105.255003>.

J W S Cook, R O Dendy, and S C Chapman. Particle-in-cell simulations of the magnetoacoustic cyclotron instability of fusion-born alpha-particles in tokamak plasmas. *Plasma Physics and Controlled Fusion*, 55(6):065003, 2013. URL <http://stacks.iop.org/0741-3335/55/i=6/a=065003>.

J. W. S. Cook, R. O. Dendy, and S. C. Chapman. Stimulated Emission of Fast Alfvén Waves within Magnetically Confined Fusion Plasmas. *Phys. Rev. Lett.*, 118:185001, May 2017. doi: 10.1103/PhysRevLett.118.185001. URL <https://link.aps.org/doi/10.1103/PhysRevLett.118.185001>.

James William Setchfield Cook. *Fusion Product Driven Lower Hybrid Electron Current in Tokamaks*. PhD thesis, The University of Warwick, 2011.

G. A. Cottrell. Identification of Minority Ion-Cyclotron Emission during Radio Frequency Heating in the JET Tokamak. *Phys. Rev. Lett.*, 84:2397–2400, Mar 2000. doi: 10.1103/PhysRevLett.84.2397. URL <https://0-link-aps-org.pugwash.lib.warwick.ac.uk/doi/10.1103/PhysRevLett.84.2397>.

G. A. Cottrell and R. O. Dendy. Superthermal Radiation from Fusion Products

in JET. *Phys. Rev. Lett.*, 60:33–36, Jan 1988. doi: 10.1103/PhysRevLett.60.33. URL <https://link.aps.org/doi/10.1103/PhysRevLett.60.33>.

G.A. Cottrell, V.P. Bhatnagar, O. Da Costa, R.O. Dendy, J. Jacquinot, K.G. McClements, D.C. McCune, M.F.F. Nave, P. Smeulders, and D.F.H. Start. Ion cyclotron emission measurements during JET deuterium-tritium experiments. *Nuclear Fusion*, 33(9):1365, 1993. URL <http://stacks.iop.org/0029-5515/33/i=9/a=I10>.

Sigmar D. J. *Proceedings of the Varenna Course on Physics of Plasmas Close to Thermonuclear Conditions*, volume I. Varenna, Commission of the European Communities, Brussels, 1979.

T. Dudok de Wit, V. V. Krasnosel'skikh, M. Dunlop, and H. Lhr. Identifying nonlinear wave interactions in plasmas using two-point measurements: A case study of Short Large Amplitude Magnetic Structures (SLAMS). *Journal of Geophysical Research: Space Physics*, 104(A8):17079–17090, 1999. doi: 10.1029/1999JA900134. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/1999JA900134>.

T. D. de Witt. *Numerical Schemes for the Analysis of Turbulence - A Tutorial (Space Plasma Simulation, Lecture Notes in Physics, vol 615) ed J. Büchner, M. Scholer, and C. T. Dum*. Berlin: Springer, 2003.

R. O. Dendy. *Plasma dynamics*. Oxford: Clarendon Press, 1990.

R. O. Dendy and K. G. McClements. Ion cyclotron wave emission at the quasiperpendicular bow shock. *Journal of Geophysical Research: Space Physics*, 98(A9):15531–15539, 1993. doi: 10.1029/93JA01386. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/93JA01386>.

R O Dendy and K G McClements. Ion cyclotron emission from fusion-born ions in large tokamak plasmas: a brief review from JET and TFTR to ITER. *Plasma Physics and Controlled Fusion*, 57(4):044002, 2015. URL <http://stacks.iop.org/0741-3335/57/i=4/a=044002>.

R. O. Dendy, C. N. LashmoreDavies, and K. F. Kam. A possible excitation mechanism for observed superthermal ion cyclotron emission from tokamak plasmas. *Physics of Fluids B: Plasma Physics*, 4(12):3996–4006, 1992. doi: 10.1063/1.860304. URL <https://doi.org/10.1063/1.860304>.

- R. O. Dendy, C. N. LashmoreDavies, and K. F. Kam. The magnetoacoustic cyclotron instability of an extended shell distribution of energetic ions. *Physics of Fluids B: Plasma Physics*, 5(7):1937–1944, 1993. doi: 10.1063/1.860781. URL <https://doi.org/10.1063/1.860781>.
- R. O. Dendy, C. N. LashmoreDavies, K. G. McClements, and G. A. Cottrell. The excitation of obliquely propagating fast alfvén waves at fusion ion cyclotron harmonics. *Physics of Plasmas*, 1(6):1918–1928, 1994a. doi: 10.1063/1.870647. URL <https://doi.org/10.1063/1.870647>.
- R. O. Dendy, K. G. McClements, C. N. LashmoreDavies, R. Majeski, and S. Cauffman. A mechanism for beamdriven excitation of ion cyclotron harmonic waves in the Tokamak Fusion Test Reactor. *Physics of Plasmas*, 1(10):3407–3413, 1994b. doi: 10.1063/1.870489. URL <https://doi.org/10.1063/1.870489>.
- R. O. Dendy, B. C. G. Reman, T. Akiyama, S. C. Chapman, J. W. S. Cook, H. Igami, S. Inagaki, K. Saito, and G. S. Yun. Proc. 44th EPS Conf. Plasma Phys. *EPS Conference Proceedings*, P5.145, 2017.
- R.O. Dendy, K.G. McClements, C.N. Lashmore-Davies, G.A. Cottrell, R. Majeski, and S. Cauffman. Ion cyclotron emission due to collective instability of fusion products and beam ions in TFTR and JET. *Nuclear Fusion*, 35(12):1733, 1995. URL <http://stacks.iop.org/0029-5515/35/i=12/a=I38>.
- R.O. Dendy, K.G. McClements, M.E. Dieckmann, and N.C. Woolsey. Energetic particles in magnetic confinement systems: synergies beyond fusion. *Nuclear Fusion*, 42(8):986, 2002. URL <http://stacks.iop.org/0029-5515/42/i=8/a=307>.
- Rudolph D’Inca. *Ion Cyclotron Emission on ASDEX Upgrade*. PhD thesis, Ludwig-Maximilians-Universität, 2014.
- H.H. Duong, W.W. Heidbrink, E.J. Strait, T.W. Petrie, R. Lee, R.A. Moyer, and J.G. Watkins. Loss of energetic beam ions during TAE instabilities. *Nuclear Fusion*, 33(5):749, 1993. URL <http://stacks.iop.org/0029-5515/33/i=5/a=I06>.
- T.Zh. Esirkepov. Exact charge conservation scheme for Particle-in-Cell simulation with an arbitrary form-factor. *Computer Physics Communications*, 135(2):144 – 153, 2001. ISSN 0010-4655. doi: [https://doi.org/10.1016/S0010-4655\(00\)00228-9](https://doi.org/10.1016/S0010-4655(00)00228-9). URL <http://www.sciencedirect.com/science/article/pii/S0010465500002289>.

- EuroFusion. The electromagnetic coil arrangement at JET. URL <https://www.euro-fusion.org/glossary/poloidal-field-coils/>.
- N. J. Fisch. Alpha power channeling using ion Bernstein waves. *Physics of Plasmas*, 2(6):2375–2380, 1995a. doi: 10.1063/1.871454. URL <https://doi.org/10.1063/1.871454>.
- N. J. Fisch. Alpha power channeling using ion Bernstein waves. *Physics of Plasmas*, 2(6):2375–2380, 1995b. doi: 10.1063/1.871454. URL <https://doi.org/10.1063/1.871454>.
- E. D. Fredrickson, M Podestá, N N Gorelenkov, D Liu, and D S Darrow. Proc. 44rd EPS Conf. Plasma Phys. *EPS Conference Proceedings*, P5.101, 2017.
- J. P Freidberg. *Ideal magnetohydrodynamics*. Plenum Press, New York, 1987.
- T. Fülöp and M. Lisak. Ion cyclotron emission from fusion products and beam ions in the Tokamak Fusion Test Reactor. *Nuclear Fusion*, 38(5):761, 1998. URL <http://stacks.iop.org/0029-5515/38/i=5/a=309>.
- T. Fülöp, Ya.I. Kolesnichenko, M. Lisak, and D. Anderson. Origin of superthermal ion cyclotron emission in tokamaks. *Nuclear Fusion*, 37(9):1281, 1997. URL <http://stacks.iop.org/0029-5515/37/i=9/a=I08>.
- T. Fülöp, M. Lisak, Ya. I. Kolesnichenko, and D. Anderson. The radial and poloidal localization of fast magnetoacoustic eigenmodes in tokamaks. *Physics of Plasmas*, 7(5):1479–1486, 2000. doi: 10.1063/1.873967. URL <https://doi.org/10.1063/1.873967>.
- P W Gingell, S C Chapman, R O Dendy, and C S Brady. Transport and evolution of ion gyro-scale plasma blobs in perpendicular magnetic fields. *Plasma Physics and Controlled Fusion*, 54(6):065005, 2012. URL <http://stacks.iop.org/0741-3335/54/i=6/a=065005>.
- P W Gingell, S C Chapman, and R O Dendy. Plasma heating by ion gyro-scale blobs in the kinetic and fluid regimes. *Plasma Physics and Controlled Fusion*, 55(5):055010, 2013. URL <http://stacks.iop.org/0741-3335/55/i=5/a=055010>.
- P W Gingell, S C Chapman, and R O Dendy. Plasma blob formation by ion kinetic kelynhelmholtz and interchange instabilities. *Plasma Physics and Controlled Fusion*, 56(3):035012, 2014. URL <http://stacks.iop.org/0741-3335/56/i=3/a=035012>.

- Peter William Gingell. *Hybrid simulations of flow bursts in magnetically confined plasmas*. PhD thesis, The University of Warwick, 2013.
- N. N. Gorelenkov. Ion cyclotron emission studies: Retrospects and prospects. *Plasma Physics Reports*, 42(5):430–439, May 2016. ISSN 1562-6938. doi: 10.1134/S1063780X16050044. URL <https://doi.org/10.1134/S1063780X16050044>.
- N. N. Gorelenkov and C. Z. Cheng. Excitation of alfvén cyclotron instability by charged fusion products in tokamaks. *Physics of Plasmas*, 2(6):1961–1971, 1995. doi: 10.1063/1.871281. URL <https://doi.org/10.1063/1.871281>.
- N.N. Gorelenkov, S.D. Pinches, and K. Toi. Energetic particle physics in fusion research in preparation for burning plasma experiments. *Nuclear Fusion*, 54(12):125001, 2014. URL <http://stacks.iop.org/0029-5515/54/i=12/a=125001>.
- G. J. Greene and TFTR. Proceedings of the 17th European Conference on Controlled Fusion and Plasma Heating, Amsterdam, Netherlands. *European Physical Society, Petit-Lancy, Switzerland*, Part IV, Vol. 14B:1540, 1990.
- K Hamamatsu, C S Chang, T Takizuka, M Azumi, T Hirayama, S Cohen, and T Tani. Numerical analysis of helium ash removal by using ICRH-driven ripple transport. *Plasma Physics and Controlled Fusion*, 40(2):255–270, feb 1998. doi: 10.1088/0741-3335/40/2/007. URL <https://doi.org/10.1088/0741-3335/40/2/007>.
- B. Hamilton, K.G. McClements, L. Fletcher, and A. Thyagaraja. Field-Guided Proton Acceleration at Reconnecting x-Points in Flares. *Solar Physics*, 214(2):339–352, Jun 2003. ISSN 1573-093X. doi: 10.1023/A:1024204928413. URL <https://doi.org/10.1023/A:1024204928413>.
- R.J. Hawryluk, D.J. Campbell, G. Janeschitz, P.R. Thomas, R. Albanese, R. Ambrosino, C. Bachmann, L. Baylor, M. Becoulet, I. Benfatto, J. Bialek, A. Boozer, A. Brooks, R. Budny, T. Casper, M. Cavinato, J.-J. Cordier, V. Chuyanov, E. Doyle, T. Evans, G. Federici, M. Fenstermacher, H. Fujieda, K. G’al, A. Garofalo, L. Garzotti, D. Gates, Y. Gribov, P. Heitzenroeder, T.C. Hender, N. Holtkamp, D. Humphreys, I. Hutchinson, K. Ioki, J. Johner, G. Johnson, Y. Kamada, A. Kavin, C. Kessel, R. Khayrutdinov, G. Kramer, A. Kukushkin, K. Lackner, I. Landman, P. Lang, Y. Liang, J. Linke, B. Lipschultz, A. Loarte, G.D. Loesser, C. Lowry, T. Luce, V. Lukash, S. Maruyama, M. Mattei, J. Menard, M. Merola, A. Mineev, N. Mitchell, E. Nardon, R. Nazikian, B. Nelson,

- C. Neumeyer, J.-K. Park, R. Pearce, R.A. Pitts, A. Polevoi, A. Portone, M. Okabayashi, P.H. Rebut, V. Riccardo, J. Roth, S. Sabbagh, G. Saibene, G. Sannazzaro, M. Schaffer, M. Shimada, A. Sen, A. Sips, C.H. Skinner, P. Snyder, R. Stambaugh, E. Strait, M. Sugihara, E. Tsitroni, J. Urano, M. Valovic, M. Wade, J. Wesley, R. White, D.G. Whyte, S. Wu, M. Wykes, and L. Zakharov. Principal physics developments evaluated in the ITER design review. *Nuclear Fusion*, 49(6):065012, 2009. URL <http://stacks.iop.org/0029-5515/49/i=6/a=065012>.
- W W Heidbrink, M E Austin, R K Fisher, M García-Munoz, G Matsunaga, G R McKee, R A Moyer, C M Muscatello, M Okabayashi, D C Pace, K Shinohara, W M Solomon, E J Strait, M A Van Zeeland, and Y B Zhu. Characterization of off-axis fishbones. *Plasma Physics and Controlled Fusion*, 53(8):085028, 2011. URL <http://stacks.iop.org/0741-3335/53/i=8/a=085028>.
- W.W. Heidbrink and G.J. Sadler. The behaviour of fast ions in tokamak experiments. *Nuclear Fusion*, 34(4):535, 1994. URL <http://stacks.iop.org/0029-5515/34/i=4/a=I07>.
- T. Hellsten and M. Laxaback. Edge localized magnetosonic eigenmodes in the ion cyclotron frequency range. *Physics of Plasmas*, 10(11):4371–4377, 2003. doi: 10.1063/1.1617315. URL <https://doi.org/10.1063/1.1617315>.
- T. Hellsten, K. Holmstrom, T. Johnson, T. Bergkvist, and M. Laxaback. On ion cyclotron emission in toroidal plasmas. *Nuclear Fusion*, 46(7):S442, 2006. URL <http://stacks.iop.org/0029-5515/46/i=7/a=S07>.
- Mark C. Herrmann and Nathaniel J. Fisch. Cooling energetic α particles in a tokamak with waves. *Phys. Rev. Lett.*, 79:1495–1498, Aug 1997. doi: 10.1103/PhysRevLett.79.1495.
- C Holland, G R Tynan, P H Diamond, R A Moyer, and M J Burin. Evidence for Reynolds-stress driven shear flows using bispectral analysis: theory and experiment. *Plasma Physics and Controlled Fusion*, 44(5A):A453, 2002. URL <http://stacks.iop.org/0741-3335/44/i=5A/a=350>.
- I. H. Hutchinson. Introduction to Plasma Physics, 2001. URL <http://silas.psfc.mit.edu/introplasma/chap2.html>.
- M. Ichimura, H. Higaki, S. Kakimoto, Y. Yamaguchi, K. Nemoto, M. Katano, M. Ishikawa, S. Moriyama, and T. Suzuki. Observation of spontaneously excited waves in the ion cyclotron frequency range on jt-60u. *Nuclear Fusion*, 48(3):035012, 2008. URL <http://stacks.iop.org/0029-5515/48/i=3/a=035012>.

- Samuel Allen Irvine. *Collective Instability and Physics of the Anomalous Doppler Resonance in Fusion Plasmas*. PhD thesis, The University of Warwick, 2018.
- S.W.A. Irvine. Kinetic dispersion solver. *GitHub repository*, 2019. URL <https://github.com/samuelirvine/Kinetic-Dispersion-Solver>.
- S-I. Itoh, K. Itoh, Y. Nagashima, and Y. Kosuga. On the Application of Cross Bispectrum and Cross Bicoherence. *Plasma and Fusion Research*, 12:1101003, 2017.
- P. Jacquet, G. BergerBy, V. Bobkov, T. Blackman, I. E. Day, F. Durodi, M. Graham, T. Hellsten, M. Laxback, M.L. Mayoral, I. Monakhov, M. Nightingale, S. E. Sharapov, M. Vrancken, and JET EFDA contributors. Parasitic signals in the receiving band of the SubHarmonic Arc Detection system on JET ICRF Antennas. *AIP Conference Proceedings*, 1406(1):17–20, 2011. doi: 10.1063/1.3664919. URL <https://aip.scitation.org/doi/abs/10.1063/1.3664919>.
- K Kamiya, N Asakura, J Boedo, T Eich, G Federici, M Fenstermacher, K Finken, A Herrmann, J Terry, A kirk, B Koch, A Loarte, R Maingi, R Maqueda, E Nardon, N Oyama, and R Sartori. Edge localized modes: recent experimental findings and related issues. *Plasma Physics and Controlled Fusion*, 49(7):S43, 2007. URL <http://stacks.iop.org/0741-3335/49/i=7/a=S03>.
- D. J. Kaup, A. Reiman, and A. Bers. Space-time evolution of non-linear three-wave interactions. I. Interaction in a homogeneous medium. *Rev. Mod. Phys.*, 51:275–309, Apr 1979. doi: 10.1103/RevModPhys.51.275. URL <https://0-link-aps-org.pugwash.lib.warwick.ac.uk/doi/10.1103/RevModPhys.51.275>.
- C.E. Kessel, D. Campbell, Y. Gribov, G. Saibene, G. Ambrosino, R.V. Budny, T. Casper, M. Cavinato, H. Fujieda, R. Hawryluk, L.D. Horton, A. Kavin, R. Kharyutdinov, F. Koechl, J. Leuer, A. Loarte, P.J. Lomas, T. Luce, V. Lukash, M. Mattei, I. Nunes, V. Parail, A. Polevoi, A. Portone, R. Sartori, A.C.C. Sips, P.R. Thomas, A. Welander, and J. Wesley. Development of ITER 15MA ELMY H-mode inductive scenario. *Nuclear Fusion*, 49(8):085034, 2009. URL <http://stacks.iop.org/0029-5515/49/i=8/a=085034>.
- Jun Young Kim, T. Rhee, Junghee Kim, S. W. Yoon, B. H. Park, M. Isobe, K. Ogawa, and W. H. Ko. Prompt loss of beam ions in kstar plasmas. *AIP Advances*, 6(10):105013, 2016. doi: 10.1063/1.4966588.

- M.H. Kim, S.G. Thatipamula, J.E. Lee, M.J. Choi, H.K. Park, T. Akiyama, and G.S. Yun. Distinct stages of radio frequency emission at the onset of pedestal collapse in KSTAR H-mode plasmas. *Nuclear Fusion*, 58(9):096034, 2018. URL <http://stacks.iop.org/0029-5515/58/i=9/a=096034>.
- Y. C. Kim and E. J. Powers. Digital Bispectral Analysis and Its Applications to Nonlinear Wave Interactions. *IEEE Transactions on Plasma Science*, 7(2):120–131, June 1979. ISSN 0093-3813. doi: 10.1109/TPS.1979.4317207.
- Y. C. Kim, J. M. Beall, E. J. Powers, and R. W. Miksad. Bispectrum and nonlinear wave coupling. *The Physics of Fluids*, 23(2):258–263, 1980. doi: 10.1063/1.862966. URL <https://aip.scitation.org/doi/abs/10.1063/1.862966>.
- H. Kimura, Y. Kusama, M. Saigusa, G.J. Kramer, K. Tobita, M. Nemoto, T. Kondoh, T. Nishitani, O. Da Costa, T. Ozeki, T. Oikawa, S. Moriyama, A. Morioka, G.Y. Fu, C.Z. Cheng, and V.I. Afanas'ev. Alfvén eigenmode and energetic particle research in JT-60U. *Nuclear Fusion*, 38(9):1303, 1998. URL <http://stacks.iop.org/0029-5515/38/i=9/a=304>.
- A. Kirk, B. Koch, R. Scannell, H. R. Wilson, G. Counsell, J. Dowling, A. Herrmann, R. Martin, and M. Walsh. Evolution of Filament Structures during Edge-Localized Modes in the MAST Tokamak. *Phys. Rev. Lett.*, 96:185001, May 2006. doi: 10.1103/PhysRevLett.96.185001. URL <https://0-link-aps-org.pugwash.lib.warwick.ac.uk/doi/10.1103/PhysRevLett.96.185001>.
- Ya.I. Kolesnichenko, T. Fülöp, M. Lisak, and D. Anderson. Localized fast magnetoacoustic eigenmodes in tokamak plasmas. *Nuclear Fusion*, 38(12):1871, 1998. URL <http://stacks.iop.org/0029-5515/38/i=12/a=311>.
- Ya.I. Kolesnichenko, M. Lisak, and D. Anderson. Superthermal radiation from tokamak plasmas caused by cyclotron magnetoacoustic instability. *Nuclear Fusion*, 40(7):1419, 2000. URL <http://stacks.iop.org/0029-5515/40/i=7/a=311>.
- V. Kravtchenko-Berejnoi, F. Lefeuvre, V. Krasnosel'skikh, and D. Lagoutte. On the use of tricoherent analysis to detect non-linear wave-wave interactions. *Signal Processing*, 42(3):291 – 309, 1995. ISSN 0165-1684. doi: [https://doi.org/10.1016/0165-1684\(94\)00136-N](https://doi.org/10.1016/0165-1684(94)00136-N). URL <http://www.sciencedirect.com/science/article/pii/016516849400136N>.
- Jaehyun Lee, Gunsu S. Yun, Minjun J. Choi, Jae-Min Kwon, Young-Mu Jeon, Woochang Lee, Neville C. Luhmann, and Hyeon K. Park. Nonlinear Interaction of Edge-Localized Modes and Turbulent Eddies in Toroidal Plasma under

$n = 1$ Magnetic Perturbation. *Phys. Rev. Lett.*, 117:075001, Aug 2016. doi: 10.1103/PhysRevLett.117.075001. URL <https://link.aps.org/doi/10.1103/PhysRevLett.117.075001>.

J Leem, G S Yun, and H K Park. Development of fast rf spectrometer system for mhd detection. *Journal of Instrumentation*, 7(01):C01042, 2012. URL <http://stacks.iop.org/1748-0221/7/i=01/a=C01042>.

A. W. Leonard. Edge-localized-modes in tokamaks. *Physics of Plasmas*, 21(9):090501, 2014. doi: 10.1063/1.4894742. URL <https://doi.org/10.1063/1.4894742>.

A Loarte, G Saibene, R Sartori, D Campbell, M Becoulet, L Horton, T Eich, A Herrmann, G Matthews, N Asakura, A Chankin, A Leonard, G Porter, G Federici, G Janeschitz, M Shimada, and M Sugihara. Characteristics of type I ELM energy and particle losses in existing devices and their extrapolation to ITER. *Plasma Physics and Controlled Fusion*, 45(9):1549, 2003. URL <http://stacks.iop.org/0741-3335/45/i=9/a=302>.

K. G. McClements and R. O. Dendy. Ion cyclotron harmonic wave generation by ring protons in space plasmas. *Journal of Geophysical Research: Space Physics*, 98(A7):11689–11700. doi: 10.1029/93JA00158. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/93JA00158>.

K. G. McClements and R. O. Dendy. Ion cyclotron harmonic wave generation by ring protons in space plasmas. *Journal of Geophysical Research: Space Physics*, 98(A7):11689–11700, 1993. doi: 10.1029/93JA00158. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/93JA00158>.

K G McClements and E D Fredrickson. Energetic particles in spherical tokamak plasmas. *Plasma Physics and Controlled Fusion*, 59(5):053001, 2017. URL <http://stacks.iop.org/0741-3335/59/i=5/a=053001>.

K. G. McClements, R. O. Dendy, and C. N. LashmoreDavies. A model for the generation of obliquely propagating ulf waves near the magnetic equator. *Journal of Geophysical Research: Space Physics*, 99(A12):23685–23693, 1994. doi: 10.1029/94JA01979. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/94JA01979>.

K. G. McClements, R. O. Dendy, C. N. LashmoreDavies, G. A. Cottrell, S. Cauffman, and R. Majeski. Interpretation of ion cyclotron emission from subAlfvnic fu-

sion products in the Tokamak Fusion Test Reactor. *Physics of Plasmas*, 3(2):543–553, 1996. doi: 10.1063/1.871881. URL <https://doi.org/10.1063/1.871881>.

K. G. McClements, C. Hunt, R. O. Dendy, and G. A. Cottrell. Ion Cyclotron Emission from JET D-T Plasmas. *Phys. Rev. Lett.*, 82:2099–2102, Mar 1999. doi: 10.1103/PhysRevLett.82.2099. URL <https://link.aps.org/doi/10.1103/PhysRevLett.82.2099>.

K.G. McClements, R. D’Inca, R.O. Dendy, L. Carbajal, S.C. Chapman, J.W.S. Cook, R.W. Harvey, W.W. Heidbrink, and S.D. Pinches. Fast particle-driven ion cyclotron emission (ICE) in tokamak plasmas and the case for an ICE diagnostic in ITER. *Nuclear Fusion*, 55(4):043013, 2015. URL <http://stacks.iop.org/0029-5515/55/i=4/a=043013>.

K.G. McClements, A. Brisset, B. Chapman, S.C. Chapman, R.O. Dendy, P. Jacquet, V.G. Kiptily, M. Mantsinen, B.C.G. Reman, and JET Contributors. Observations and modelling of ion cyclotron emission observed in JET plasmas using a sub-harmonic arc detection system during ion cyclotron resonance heating. *Nuclear Fusion*, 58(9):096020, 2018. URL <http://stacks.iop.org/0029-5515/58/i=9/a=096020>.

R. A. Moyer, G. R. Tynan, C. Holland, and M. J. Burin. Increased Nonlinear Coupling between Turbulence and Low-Frequency Fluctuations at the L-H Transition. *Phys. Rev. Lett.*, 87:135001, Sep 2001. doi: 10.1103/PhysRevLett.87.135001. URL <https://link.aps.org/doi/10.1103/PhysRevLett.87.135001>.

R. Ochoukov, V. Bobkov, H. Faugel, H. Fngfelder, and J.-M. Noterdaeme. A new B-dot probe-based diagnostic for amplitude, polarization, and wavenumber measurements of ion cyclotron range-of-frequency fields on ASDEX Upgrade. *Review of Scientific Instruments*, 86(11):115112, 2015. doi: 10.1063/1.4935833. URL <https://doi.org/10.1063/1.4935833>.

R. Ochoukov, V. Bobkov, B. Chapman, R. Dendy, M. Dunne, H. Faugel, M. Garca-Muoz, B. Geiger, P. Hennequin, K. G. McClements, D. Moseev, S. Nielsen, J. Rasmussen, P. Schneider, M. Weiland, and J.-M. Noterdaeme. Observations of core ion cyclotron emission on ASDEX Upgrade tokamak. *Review of Scientific Instruments*, 89(10):10J101, 2018. doi: 10.1063/1.5035180. URL <https://doi.org/10.1063/1.5035180>.

R. Ochoukov, R. Bilato, V. Bobkov, B. Chapman, S.C. Chapman, R.O. Dendy, M. Dunne, H. Faugel, M. Garca-Muoz, B. Geiger, A. Kallenbach, A. Kappatou,

- K.G. McClements, D. Moseev, S. Nielsen, J. Rasmussen, P. Schneider, M. Weiland, J.-M. Noterdaeme, ASDEX Upgrade Team, and EUROfusion MST1 Team. Core plasma ion cyclotron emission driven by fusion-born ions. *Nuclear Fusion*, 59(1):014001, 2019. URL <http://stacks.iop.org/0029-5515/59/i=1/a=014001>.
- D.C. Pace, W.W. Heidbrink, R.I. Pinsky, M.A. Van Zeeland, and Y.B. Zhu. Proc. 43rd EPS Conf. Plasma Phys. O2.101, 2016.
- J. L. Posch, M. J. Engebretson, C. N. Olson, S. A. Thaller, A. W. Breneman, J. R. Wygant, S. A. Boardsen, C. A. Kletzing, C. W. Smith, and G. D. Reeves. Lowharmonic magnetosonic waves observed by the Van Allen Probes. *Journal of Geophysical Research: Space Physics*, 120(8):6230–6257, 2015. doi: 10.1002/2015JA021179. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2015JA021179>.
- V. L. Rekaa, S. C. Chapman, and R. O. Dendy. Ion Pre-acceleration in Fully Self-consistent Particle-in-cell Simulations of Supercritical Perpendicular Reforming Shocks in Multiple Ion Species Plasmas. *The Astrophysical Journal*, 791(1):26, 2014. URL <http://stacks.iop.org/0004-637X/791/i=1/a=26>.
- B. C. G. Reman, R. O. Dendy, T. Akiyama, S. C. Chapman, J. W. S. Cook, H. Igami, S. Inagaki, K. Saito, and G. S. Yun. Proc. 43rd EPS Conf. Plasma Phys. *EPS Conference Proceedings*, P2.041, 2016.
- K. Saito, H. Kasahara, T. Seki, R. Kumazawa, T. Mutoh, T. Watanabe, F. Shimpo, G. Nomura, M. Osakabe, M. Ichimura, H. Higaki, and A. Komori. Measurement of ion cyclotron emissions by use of ICRF heating antennas in LHD. *Fusion Engineering and Design*, 84(7):1676 – 1679, 2009. ISSN 0920-3796. doi: <https://doi.org/10.1016/j.fusengdes.2008.12.053>. URL <http://www.sciencedirect.com/science/article/pii/S0920379608005073>. Proceeding of the 25th Symposium on Fusion Technology.
- Kenji Saito, Ryuhei Kumazawa, Tetsuo Seki, Hiroshi Kasahara, Goro Nomura, Fujio Shimpo, Hiroe Igami, Mitsutaka Isobe, Kunihiro Ogawa, Kazuo Toi, Masaki Osakabe, Masaki Nishiura, Tsuguhiro Watanabe, Satoshi Yamamoto, Makoto Ichimura, Takashi Mutoh, and LHD Experiment Group. Measurement of Ion Cyclotron Emissions by Using High-Frequency Magnetic Probes in the LHD. *Plasma Science and Technology*, 15(3):209, 2013. URL <http://stacks.iop.org/1009-0630/15/i=3/a=03>.

- S Sato, M Ichimura, Y Yamaguchi, M Katano, Y Imai, T Murakami, Y Miyake, T Yokoyama, S Moriyama, T Kobayashi, A Kojima, K Shinohara, Y Sakamoto, T Watanabe, H Hojo, and T Imai. Observation of Ion Cyclotron Emission Owing to DD Fusion Product H Ions in JT-60U. *Plasma and Fusion Research*, 5:S2067, 2010.
- O. Sauter and S.Yu. Medvedev. Tokamak coordinate conventions: COCOS. *Computer Physics Communications*, 184(2):293 – 302, 2013. ISSN 0010-4655. doi: <https://doi.org/10.1016/j.cpc.2012.09.010>. URL <http://www.sciencedirect.com/science/article/pii/S0010465512002962>.
- P. Schild, G.A. Cottrell, and R.O. Dendy. Sawtooth oscillations in ion cyclotron emission from JET. *Nuclear Fusion*, 29(5):834, 1989. URL <http://stacks.iop.org/0029-5515/29/i=5/a=013>.
- O. Schmitz, K. Ida, M. Kobayashi, A. Bader, S. Brezinsek, T.E. Evans, H. Funaba, M. Goto, O. Mitarai, T. Morisaki, G. Motojima, Y. Nakamura, Y. Narushima, D. Nicolai, U. Samm, H. Tanaka, H. Yamada, M. Yoshinuma, and Y. Xu and. Enhancement of helium exhaust by resonant magnetic perturbation fields at LHD and TEXTOR. *Nuclear Fusion*, 56(10):106011, aug 2016. doi: 10.1088/0029-5515/56/10/106011. URL <https://doi.org/10.1088/0029-5515/56/10/106011>.
- Alexander Shalashov, E V Suvorov, L V Lubyako, H Maassberg, and the Team. Nbi-driven ion cyclotron instabilities at the w7-as stellarator. 45:395, 03 2003.
- H. Smith, T. Flp, M. Lisak, and D. Anderson. Localization of compressional alfvén eigenmodes in spherical tori. *Physics of Plasmas*, 10(5):1437–1442, 2003. doi: 10.1063/1.1566441. URL <https://doi.org/10.1063/1.1566441>.
- H M Smith and E Verwichte. Compressional alfvén eigenmode structure in spherical tokamaks. *Plasma Physics and Controlled Fusion*, 51(7):075001, 2009. URL <http://stacks.iop.org/0741-3335/51/i=7/a=075001>.
- L. S. Solov'ev. Magnetoacoustic cyclotron instability in a thermonuclear plasma. *Sov. Phys.-JETP*, 26(6):400, 1968.
- T. H. Stix. *Waves in Plasmas*, pages 3–9. Springer-Verlag New York, Inc., 1992.
- Shuhei Sumida, Kouji Shinohara, Ryuya Ikezoe, Makoto Ichimura, Mizuki Sakamoto, Mafumi Hirata, and Shunsuke Ide. Comparison of Dispersion Model of Magneto-Acoustic Cyclotron Instability with Experimental Observation of ^3He

- Ion Cyclotron Emission on JT-60U. *Journal of the Physical Society of Japan*, 86(12):124501, 2017. doi: 10.7566/JPSJ.86.124501. URL <https://doi.org/10.7566/JPSJ.86.124501>.
- Shuhei Sumida, Kouji Shinohara, Ryuya Ikezoe, Makoto Ichimura, Mizuki Sakamoto, Mafumi Hirata, and Shunsuke Ide. Characteristics of fast ^3He ion velocity distribution exciting ion cyclotron emission on JT-60U. *Plasma Physics and Controlled Fusion*, 2018. URL <http://iopscience.iop.org/10.1088/1361-6587/aaf184>.
- R.D. Sydora. Low-noise electromagnetic and relativistic particle-in-cell plasma simulation models. *Journal of Computational and Applied Mathematics*, 109(1):243 – 259, 1999. ISSN 0377-0427. doi: [https://doi.org/10.1016/S0377-0427\(99\)00161-2](https://doi.org/10.1016/S0377-0427(99)00161-2). URL <http://www.sciencedirect.com/science/article/pii/S0377042799001612>.
- Toshio Terasawa, Masahiro Hoshino, Jun-Ichi Sakai, and Tohru Hada. Decay instability of finite-amplitude circularly polarized alfvén waves: A numerical simulation of stimulated brillouin scattering. *Journal of Geophysical Research: Space Physics*, 91(A4):4171–4187, 1986. doi: 10.1029/JA091iA04p04171. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/JA091iA04p04171>.
- Shekar G Thatipamula, G S Yun, J Leem, H K Park, K W Kim, T Akiyama, and S G Lee. Dynamic spectra of radio frequency bursts associated with edge-localized modes. *Plasma Physics and Controlled Fusion*, 58(6):065003, 2016. URL <http://stacks.iop.org/0741-3335/58/i=6/a=065003>.
- K. E. Thome, D. C. Pace, R. I. Pinsky, O. Meneghini, C. A. del Castillo, and Y. Zhu. Radio frequency measurements of energetic-particle-driven emission using the ion cyclotron emission diagnostic on the DIII-D tokamak. *Review of Scientific Instruments*, 89(10):10I102, 2018. doi: 10.1063/1.5035561. URL <https://doi.org/10.1063/1.5035561>.
- J.-L. Vay. Simulation of beams or plasmas crossing at relativistic velocity. *Physics of Plasmas*, 15(5):056701, 2008. doi: 10.1063/1.2837054. URL <https://doi.org/10.1063/1.2837054>.
- A. L. Verdon, Iver H. Cairns, D. B. Melrose, and P. A. Robinson. Warm electromagnetic lower hybrid wave dispersion relation. *Physics of Plasmas*, 16(5):052105, 2009. doi: 10.1063/1.3132628. URL <https://doi.org/10.1063/1.3132628>.

- J Villasenor and O Buneman. Rigorous charge conservation for local electromagnetic field solvers. *Computer Physics Communications*, 69:306–316, 1992. doi: 10.1016/0010-4655(92)90169-Y.
- F. Wagner, G. Becker, K. Behringer, D. Campbell, A. Eberhagen, W. Engelhardt, G. Fussmann, O. Gehre, J. Gernhardt, G. v. Gierke, G. Haas, M. Huang, F. Karger, M. Keilhacker, O. Klüber, M. Kornherr, K. Lackner, G. Lisitano, G. G. Lister, H. M. Mayer, D. Meisel, E. R. Müller, H. Murmann, H. Niedermeyer, W. Poschenrieder, H. Rapp, H. Röhr, F. Schneider, G. Siller, E. Speth, A. Stäbler, K. H. Steuer, G. Venus, O. Vollmer, and Z. Yü. Regime of Improved Confinement and High Beta in Neutral-Beam-Heated Divertor Discharges of the ASDEX Tokamak. *Phys. Rev. Lett.*, 49:1408–1412, Nov 1982. doi: 10.1103/PhysRevLett.49.1408. URL <https://link.aps.org/doi/10.1103/PhysRevLett.49.1408>.
- J. Wesson. *Tokamaks 3rd edition*, pages 15–18. Calarendon Press - Oxford, 2004.
- D. Winske, L. Yin, N. Omid, H. Karimabadi, and K. Quest. *Hybrid Simulation Codes: Past, Present and Future A Tutorial (Space Plasma Simulation, Lecture Notes in Physics, vol 615) ed J. Büchner, M. Scholer, and C. T. Dum*. Berlin: Springer, 2003.
- Takuma Yamada, Sanae-I. Itoh, Takashi Maruta, Naohiro Kasuya, Yoshihiko Nagashima, Shunjiro Shinohara, Kenichiro Terasaka, Masatoshi Yagi, Shigeru Inagaki, Yoshinobu Kawai, Akihide Fujisawa, and Kimitaka Itoh. Anatomy of plasma turbulence. *Nature Physics*, 4:721–725, 2008.
- Kane Yee. Numerical solution of initial boundary value problems involving maxwell’s equations in isotropic media. *IEEE Transactions on Antennas and Propagation*, 14(3):302–307, May 1966. ISSN 0018-926X. doi: 10.1109/TAP.1966.1138693.
- G. S. Yun, W. Lee, M. J. Choi, J. Lee, H. K. Park, B. Tobias, C. W. Domier, N. C. Luhmann, A. J. H. Donné, and J. H. Lee. Two-Dimensional Visualization of Growth and Burst of the Edge-Localized Filaments in KSTAR H-Mode Plasmas. *Phys. Rev. Lett.*, 107:045004, Jul 2011. doi: 10.1103/PhysRevLett.107.045004. URL <https://link.aps.org/doi/10.1103/PhysRevLett.107.045004>.
- H Zohm. Edge localized modes (ELMs). *Plasma Physics and Controlled Fusion*, 38(2):105, 1996. URL <http://stacks.iop.org/0741-3335/38/i=2/a=001>.