

Modelling runaway electron generation in tokamaks

LORENZO VOTTA

Academic Dissertation which, with due permission of the KTH Royal Institute of Technology, is submitted for public defence for the Degree of Licentiate of Engineering on Friday the 12th of June 2026, at 10:00 a.m. in H1, Teknikringen 33, Stockholm.

Licentiate Thesis in Electrical Engineering
KTH Royal Institute of Technology
Stockholm, Sweden 2026

© Lorenzo Votta

Cover page photo: Lorenzo Votta

TRITA-EECS-AVL-2026:59

ISBN 978-91-8106-637-1

Printed by: Universitetservice US-AB, Sweden 2026

Abstract

Tokamak disruptions can convert a large fraction of the plasma current into a beam of relativistic runaway electrons. In a reactor-scale device such as ITER, a runaway electron beam could carry several megaamperes and, if left uncontrolled, could cause severe damage to plasma-facing components. Predicting whether a given disruption scenario leads to a dangerous runaway beam, and designing injection schemes that prevent it, requires models that capture the interplay between material injection, rapid plasma cooling, electric field evolution, and the various mechanisms by which runaway electrons are born, multiply, and are lost. This thesis addresses runaway electron physics from seed formation to disruption mitigation through numerical modelling.

A synthetic electron cyclotron emission (ECE) framework is developed and applied to vertical ECE measurements on the TCV tokamak, combining Fokker-Planck calculations of the electron distribution function with ray tracing and radiative transfer. The analysis demonstrates that vertical ECE can resolve the energy-dependent dynamics of suprathermal electrons in the 20–100 keV range, providing constraints on the nascent runaway seed that are difficult to obtain with conventional diagnostics.

The disruption simulation framework DREAM is then extended with several physics models relevant to ITER: runaway electron losses from vertical plasma displacement, cross-field drift of pellet ablation material, stochasticity driven current-profile relaxation, and an updated Compton scattering source for the ITER first wall. These are applied to a systematic study of shattered pellet injection scenarios in ITER showing that avoiding a multi-megaampere runaway beam depends sensitively on the thermal quench timescale, the injected material composition, and the competition between runaway multiplication and scrape-off losses. Finally, a viable theoretical pathway that limits the runaway current to tolerable levels even in the presence of nuclear runaway sources is identified.

Keywords: Nuclear fusion, Tokamak, Runaway electrons, Disruptions

Sammanfattning

Vid avbrott i en tokamak kan en stor del av plasmaströmmen omvandlas till en stråle av relativistiska skenande elektroner. I en reaktor i ITER-skala skulle en sådan stråle kunna bära flera megaampere ström och orsaka allvarliga skador på komponenter nära plasmat om den inte kontrolleras. För att förutsäga om ett visst avbrott leder till en farlig stråle av skenande elektroner och utforma injektionsmetoder som kan förhindra detta krävs modeller som fångar samspelen mellan materialinjektion, snabb plasmakylning, det elektriska fältets utveckling samt de olika mekanismer genom vilka skenande elektroner uppstår, multipliceras och förloras. Denna avhandling behandlar fysiken bakom skenande elektroner från deras bildande till mitigering av avbrott, med hjälp av numerisk modellering.

Ett syntetiskt ramverk för elektron-cyklotron-strålning (ECE) har utvecklats och tillämpats på vertikala ECE-mätningar vid experimentet TCV. Fokker-Planck-beräkningar av elektronernas fördelningsfunktion kombineras med stråloptik och strålningstransport. Analysen visar att vertikal ECE kan upplösa den energiberoende dynamiken hos supratherma elektroner i energintervallet 20–100 keV, vilket ger information om den begynnande populationen av skenande elektroner som annars är svår att erhålla med konventionell diagnostik.

Det integrerade ramverket för avbrottssimuleringar DREAM har därefter utökats med flera fysikmodeller relevanta för ITER: förluster av skenande elektroner till följd av vertikal plasmaförskjutning, tvärfältsdrift av material från pelletablation, stokasticitetsdriven relaxation av strömprofilen samt en uppdaterad källterm för Compton-spridning motsvarande ITER:s primära vägg. Dessa modeller har tillämpats i en systematisk studie av scenarier med injektion av pellet-splitter i ITER, vilken visar att möjligheten att undvika flera megaampere av skenande elektronström beror starkt på tidsskalan för värmeförlusten, sammansättningen hos det injicerade materialet samt balansen mellan multiplikation av skenande elektroner och hur snabbt de skrapas av mot väggen. Slutligen identifieras en teoretiskt möjlig väg framåt för ITER som begränsar strömmen som bärs av de skenande elektronerna till tolerabla nivåer, även i de fall då skenande elektroner kan uppstå genom nukleära processer.

Publications

- I *Experimental and numerical investigation of suprathermal electron dynamics using vertical electron cyclotron emission,*
L. Votta, M. Hoppe, J. Decker, E. Devlaminck, A. S. Tema Biwolé, L. Porte, J. Cazabonne, Y. Savoye-Peysson and the TCV Team
Plasma Physics and Controlled Fusion, vol. **68**, p. 015029, (2026)
- II *Runaway electron generation in ITER mitigated disruptions with improved physics models,*
L. Votta, F. J. Artola, E. Nardon, O. Vallhagen, M. Hoppe
Submitted to Nuclear Fusion (2026)
- III *Reduced modelling of scrape-off losses of runaway electrons during tokamak disruptions,*
O. Vallhagen, L. Hanebring, T. Fülöp, M. Hoppe, **L. Votta**, I. Pusztai
Journal of Plasma Physics, vol. **91**, no. 3, E78, (2025)
- IV *Simulation of shattered pellet injections with plasmoid drifts in ASDEX Upgrade and ITER,*
O. Vallhagen, L. Antonsson, P. Halldestam, G. Papp, P. Heinrich, A. Patel, M. Hoppe, **L. Votta**, the ASDEX Upgrade Team, the EUROfusion Tokamak Exploitation Team
Plasma Physics and Controlled Fusion, vol. **67**, no. 10, p. 105034, (2025)

Statement of contribution

- Paper I I designed and implemented the YODA synthetic diagnostic and its coupling to the LUKE kinetic solver and the C3PO ray-tracing code. I validated the framework through a systematic benchmark against the SPECE synthetic diagnostic, and carried out all simulations presented in the paper. I performed the data analysis and was leading the physical interpretation of the results. I wrote the manuscript together with the co-authors.
- Paper II I developed and implemented the disruption simulation framework used to carry out systematic studies of SPI-triggered disruptions in ITER with DREAM, and generalised it for application to other tokamaks. I implemented the halo heat loss term in the cold electron energy balance equation, and made substantial contributions to the development of the hyperresistive diffusion model. I designed and executed all simulations presented in the paper, performed the data analysis, and was leading the modelling choices and the physical interpretation of the results. I wrote the manuscript together with the co-authors.
- Paper III I contributed to the development and implementation of the runaway electron scrape-off term in DREAM, with a particular focus on its numerical integration into the simulation framework. I performed the benchmark of this model against 2D JOREK simulations, and made significant contributions to the analysis and interpretation of the results.
- Paper IV I contributed to the validation of the plasmoid drift model implemented in DREAM, and made substantial contributions to the analysis and physical interpretation of the ITER simulations of plasmoid drift dynamics.

Related publications, not included in the thesis

- V *Expulsion of runaway electrons using ECRH in the TCV tokamak*
J. Decker, M. Hoppe, U. Sheikh, B.P. Duval, G. Papp, L. Simons,
T. Wijkamp, J. Cazabonne, S. Coda, E. Devlaminck, O. Ficker, R.
Hellinga, U. Kumar, Y. Savoye-Peysson, L. Porte, C. Reux, C. Som-
mariva, A. Tema Biwolé, B. Vincent, **L. Votta** and the TCV Team
Nuclear Fusion, vol. **64**, no 10 (2024)
- VI *Cross-calibration and first vertical ECE measurement of electron en-
ergy distribution in the TCV tokamak*
A. Tema Biwole, L. Porte, A. Fasoli, L. Figini, J. Decker, M. Hoppe,
J. Cazabonne, **L. Votta**, A. Simonetto, S. Coda and the TCV Team
Plasma Physics and Controlled Fusion, vol. **66**, no 12 (2024)

Conference contributions

- VII *Runaway electron dynamics in the Tokamak à Configuration Variable*
M. Hoppe, J. Decker, U. Sheikh, T. Wijkamp, J. Cazabonne, B. P. Duval, G. Papp, A. Perek, Y. Peysson, L. Simons, A. Tema-Biwole, B. Vincent, **L. Votta** and the TCV team *Proceedings of the 49th EPS Conference on Plasma Physics, Bordeaux (2023)*
- VIII *Experimental and numerical investigations of suprathermal electron dynamics in TCV using electron cyclotron emission*
L. Votta, M. Hoppe, J. Decker, A. Tema-Biwole, L. Porte, J. Cazabonne, S. Coda, Y.S. Peysson
Proceedings of the 50th EPS Conference on Plasma Physics, Salamanca (2024)
- IX *Turbulence-inclusive modelling of electron cyclotron wave-plasma dynamics in tokamaks*
E. Devlaminck, J. Decker, S. Coda, L. Porte, J. Cazabonne, O. Maj, E. Poli, **L. Votta**, Y. Peysson and the TCV team
Proceedings of the 51st EPS Conference on Plasma Physics, Vilnius (2025)

Acknowledgement

Writing acknowledgements is always a difficult task. Few combinations of words can properly express how fortunate I feel to be doing this work, in this place, and with the people who have made it possible. Still, I will try my best.

First, I would like to thank my principal supervisor, Prof. Mathias Hoppe, for his guidance throughout these years. He is, in many ways, the reason I am here. He trusted me from the very beginning, when he supervised my Master's thesis at EPFL in Lausanne, and that trust has continued through my licentiate work and into the future of my PhD. His supervision has been essential not only for the results presented here, but also for the way I am learning to think as a researcher.

I also thank my co-supervisor, Prof. Thomas Jonsson, for many valuable physics discussions, for his passionate and careful perspective, and for the support he has provided during this first half of my doctoral studies.

I am grateful to Prof. Lorenzo Frassinetti for taking the time to carefully read this thesis, and for the years I spent teaching in his *Vektoranalys* course. That experience has been an important part of my time at KTH, both scientifically and personally.

Several collaborators across Europe have contributed, directly or indirectly, to the work presented in this thesis. Among them, I am particularly grateful to Joan Decker, Javier Artola, Eric Nardon, Oskar Vallhagen, and my friends and colleagues in the DREAM team. Their expertise, discussions, and feedback have been invaluable in helping me place my work within a broader scientific context.

I also thank my colleagues at KTH, and especially the old and new members of the “cool office”: Laura, Hampus, Björn, Lukas, Linn, Arnaud, Olle and Fotios. This job is rarely only about the research itself; the people with whom one shares offices, coffee breaks, frustrations, jokes, and eventually a whiteboard full of cows, matter more than one expects.

I also want to thank my friends, in Stockholm and elsewhere, those who have been part of my life for many years and those I have met along the way. I will not list names here, to avoid making these acknowledgments unnecessarily long. If you are reading this, you know how much you matter to me.

To Veronica, I do not even know where to begin. Thank you for your patience, support, presence and love. You have been there through the ordinary days and the difficult ones, and that has made a difference much larger than what can be properly expressed here.

Finally, I thank my family. *Prima di tutti, Nonna Rosetta: grazie per il tuo amore incondizionato, che porto sempre con me ovunque io sia.* I thank my parents, Mafalda and Lello, and my brother Giorgio, who have supported me from the beginning, each in their own way, with trust, affection, and the quiet certainty that I would find my path. Much of what I have done, and much of what I will do, rests on that foundation.

Lorenzo Votta
Stockholm, May 2026

Acronyms

List of commonly used acronyms:

ADAS	Atomic Data and Analysis Structure
AUG	ASDEX Upgrade
CQ	Current Quench
DMS	Disruption Mitigation System
DREAM	Disruption Runaway Electron Analysis Model
EC	Electron Cyclotron
ECCD	Electron Cyclotron Current Drive
ECE	Electron Cyclotron Emission
ECRH	Electron Cyclotron Resonance Heating
HXR	Hard X-Ray
JET	Joint European Torus
KSTAR	Korea Superconducting Tokamak Advanced Research
LCFS	Last Closed Flux Surface
LOS	Line of Sight
MGI	Massive Gas Injection
MHD	Magnetohydrodynamics
MMI	Massive Material Injection
NGS	Neutral Gas Shielding
NTM	Neoclassical Tearing Mode
RE	Runaway Electron
SPI	Shattered Pellet Injection
STEP	Spherical Tokamak for Energy Production
TCV	Tokamak à Configuration Variable
TQ	Thermal Quench
VDE	Vertical Displacement Event
VECE	Vertical Electron Cyclotron Emission
WKB	Wentzel–Kramers–Brillouin

Contents

Publications	iii
Acknowledgement	vii
Acronyms	ix
Contents	1
1 Introduction	3
1.1 Global energy demand	3
1.2 Nuclear fusion	5
1.3 Magnetic confinement and the tokamak concept	7
1.4 Disruptions and runaway electrons	9
1.4.1 Disruptions	9
1.4.2 Runaway electrons	11
1.5 Thesis outline	13
2 Runaway electrons	15
2.1 The kinetic equation	15
2.2 Runaway generation	17
2.2.1 Primary mechanisms	19
2.2.2 Secondary mechanism – avalanche generation	23
2.3 Runaway electron losses	25
2.3.1 Transport losses	25
2.3.2 Scrape-off losses	26
3 Disruption mitigation	29
3.1 Massive material injection	30
3.1.1 Massive gas injection	30
3.1.2 Shattered pellet injection	31

3.2	Disruption modelling	33
3.2.1	Material injection	33
3.2.2	Material ablation and deposition	34
3.2.3	Density and charge-state evolution	37
3.2.4	Energy balance and thermal quench	38
3.2.5	Electric field and currents	41
3.2.6	Runaway electrons	42
3.3	Numerical tools	44
4	Suprathermal electron cyclotron emission	45
4.1	Vertical Electron Cyclotron Emission	47
4.2	Suprathermal ECE modelling	48
4.2.1	Wave propagation	49
4.2.2	Radiation transport	49
4.2.3	Emission and absorption coefficients	50
4.3	Numerical tools	52
4.3.1	LUKE: Fokker–Planck solver	52
4.3.2	C3PO: ray tracing	53
4.3.3	YODA: synthetic ECE diagnostic	53
5	Summary and outlook	55
5.1	Paper I – Experimental and numerical investigation of suprathermal electron dynamics using vertical electron cyclotron emission	55
5.2	Paper II – Runaway electron generation in ITER mitigated disruptions with improved physics models	56
5.3	Paper III – Reduced modelling of scrape-off losses of runaway electrons during tokamak disruptions	57
5.4	Paper IV – Simulation of shattered pellet injections with plasmoid drifts in ASDEX Upgrade and ITER	58
5.5	Outlook	58
	References	61

Chapter 1

Introduction

For decades, the prospect of abundant and low carbon energy from controlled nuclear fusion on Earth has been an ambitious goal. Despite sustained scientific and technological progress, substantial challenges remain. Among the most concerning is the generation of large and energetic beams of relativistic electrons during tokamak disruptions. These are commonly referred to as runaway electrons and are the focus of this thesis.

To motivate their study, this chapter first clarifies the role of fusion as an energy source and introduces the basic principles of magnetic confinement, with emphasis on the tokamak concept and briefly introduces the concepts of disruptions and runaway electrons.

1.1 Global energy demand

Global energy demand continues to rise to unprecedented levels. In 2023, worldwide primary energy consumption reached a record high of about 620 exajoules (EJ), having grown at an average rate of roughly 1.5% per year over the past decade [1]. This upward trend is driven by a growing population and rising standards of living: advances in energy efficiency have been unable to offset increased demand from population growth and higher per-capita consumption as developing regions industrialize [2]. In a current-policies scenario, global primary energy use in 2050 is projected to be 16–57% higher than in 2022 [2], continuing the historical growth trajectory. Providing secure, sufficient energy to a larger and more affluent world is thus a central challenge for the coming decades.

Equally important is doing so sustainably. Fossil fuels still supply more than 80% of world primary energy today, and are responsible for the major-

ity of anthropogenic greenhouse gas emissions [1]. The resulting climate impacts, together with air pollution and other environmental pressures, imply that a major energy transition is needed to meet global energy needs while eliminating net carbon emissions.

Recent years have seen rapid expansion of renewable energy sources, but current renewables alone face limitations in meeting global demand at scale. Solar and wind power are variable by their intermittent nature, requiring energy storage or backup generation to ensure a continuous supply. They also have relatively low energy density, implying large land, or offshore area requirements per unit of energy. For example, supplying today's energy consumption using only solar photovoltaics has been estimated to require on the order of a 550 km by 550 km area of panels (roughly 0.2% of Earth's land area), illustrating the immense scale of infrastructure needed. Moreover, not all energy end-uses are easily electrified, and integration challenges and reliability issues grow as variable renewables approach a dominant share of the supply [3]. Continued progress in batteries, grid expansion, and demand management can enable very high renewable penetrations, but it is likely that additional low-carbon, dispatchable energy sources will be needed to fully decarbonize the system.

Against this backdrop, nuclear energy provides a high-density, low-carbon option. Current fission reactors supply low-emission electricity in many countries, yet concerns about safety, waste and proliferation still limit their broader acceptance. Fusion, by contrast, aims to deliver the benefits of a nuclear energy source with inherently favourable safety characteristics. Its development is driven not only by climate considerations but also by the need for a sustainable, large-scale power source as global demand grows through mid-century. Fusion offers an energy density of $\sim 3.4 \times 10^{14}$ J/kg for the deuterium–tritium reaction [4], roughly seven orders of magnitude above chemical fuels, and relies on fuel resources (deuterium from water and lithium for tritium breeding) that are effectively inexhaustible on human timescales [5]. A reactor would produce electricity with negligible greenhouse-gas emissions and without long-lived radioactive waste [5], while delivering steady, dispatchable power. Unlike fission, it involves no risk of runaway chain reactions or large-scale melt-down and poses minimal proliferation concerns. Most importantly from a systems perspective, the enormous energy density of fusion decouples fuel availability from the scale of infrastructure required to deliver a given power output: tiny quantities of widely accessible fuel can sustain very

large power outputs. In contrast to fossil fuels and many renewables, fusion's long-term potential is not constrained by resource limits or land use, but by our ability to resolve the remaining physics and engineering challenges.

1.2 Nuclear fusion

At the heart of the fusion energy concept lies a nuclear process in which two light nuclei fuse to form a heavier one, releasing energy in the process. To understand on general grounds why fusion of light elements can be an energy source, it is useful to introduce the binding energy of a nucleus. For a nucleus in its ground state, the binding energy (B.E.) is defined as

$$\text{B.E.} = Zm_{\text{p}}c^2 + (A - Z)m_{\text{n}}c^2 - U$$

where Z is the atomic number, A is the mass number, m_{p} is the proton mass, m_{n} is the neutron mass, c is the speed of light in vacuum, while U is the total energy (rest energy) of the nucleus.




Figure 1.1: Binding energy (B.E.) per nucleon as a function of the mass number. Red markers highlight isotopes of special significance.

Figure 1.1 shows the binding energy per nucleon as a function of mass number. Nuclei lighter than iron have, on average, a lower binding energy per nucleon than medium-mass nuclei. When two light nuclei fuse to form a heavier one, the product typically lies higher on this curve, and the increase in total binding energy is released as kinetic energy of the reaction products. Conversely, very heavy nuclei can release energy by

splitting (fission) into fragments closer to iron. In both cases, energy can be extracted by exploiting the increase in binding energy per nucleon between reactants and products.

Energy-producing fusion is most favourable for the lightest nuclides for two independent reasons. First, light reactants lie on the rising side of the binding-energy-per-nucleon curve, so the fusion product is more tightly bound and the corresponding increase in total binding energy is released as kinetic energy. Second, charged nuclei experience a Coulomb repulsion whose potential barrier scales with the charge product $Z_1 Z_2$. The quantum-tunnelling probability through this barrier decreases approximately as $\exp(-\sqrt{E_G/E})$ where E is the centre-of-mass energy and $E_G \propto \mu(Z_1 Z_2)^2$ is the Gamow energy, with μ the reduced mass. Hydrogenic isotopes minimise both $Z_1 Z_2$ and μ , leading to comparatively large tunnelling probabilities at temperatures that are conceivably attainable in the laboratory [6].




Figure 1.2: Maxwellian-averaged fusion reactivity $\langle\sigma v\rangle$ as a function of temperature T for the D–D (black), D–T (green), and D– ^3He (red) reactions. Both axes are logarithmic; the D–T channel attains the largest reactivity, peaking near $T \sim 100$ keV.

In a thermonuclear plasma, fusion reaction rates are characterised by the Maxwellian-averaged reactivity $\langle\sigma v\rangle$, which depends on the plasma temperature. Figure 1.2 shows $\langle\sigma v\rangle$ as a function of temperature T for the most relevant light-element fuel pairs: D–T, D–D and D– ^3He . The D–T channel attains the largest reactivity and does so at the lowest tem-

peratures (tens of keV; the peak of $\langle\sigma v\rangle$ lies around $T \sim 100$ keV). Consequently, the main fuel candidates of practical interest for terrestrial fusion are the light–element pairs D–T, D–D and D– ^3He , with D–T being the most favourable. Specifically, the D–T reaction



partitions about 20% of the released energy to a charged α particle and about 80% to a neutron. The α particles are confined by the magnetic field and can transfer their energy back to the plasma, providing self-heating that is essential for ignition. The high D–T reactivity lowers the confinement requirements relative to alternative fuels, while the neutrons carry most of the energy to the surrounding blanket, where it can be converted to heat and, in a reactor, used to breed tritium from lithium. These features, together with deuterium’s abundance and the possibility of tritium breeding, make the D–T reaction the near-term choice to produce energy from fusion on Earth.

1.3 Magnetic confinement and the tokamak concept

To exploit the D–T fusion reaction discussed in Section 1.2 as an energy source, the fuel must be heated to temperatures of order 10–20 keV. Under such conditions, the gas is fully ionised and exists in the plasma state: a (quasi-)neutral mixture of charged particles whose dynamics is dominated by collective electromagnetic interactions rather than by binary collisions.

A simple set of criteria to define a *plasma* emphasises three basic features. First, charge perturbations are screened over the Debye length

$$\lambda_D = \sqrt{\frac{\varepsilon_0 k_B T_e}{n_e e^2}}, \quad (1.1)$$

meaning that electric fields are shielded on scales $\gtrsim \lambda_D$; second, the plasma parameter $\Lambda \equiv n\lambda_D^3 \gg 1$, ensures that many particles reside within a Debye sphere and that the system is weakly coupled. Third, the characteristic collective time scale is set by the electron plasma frequency

$$\omega_{pe} = \sqrt{\frac{n_e e^2}{\varepsilon_0 m_e}}, \quad (1.2)$$

which is typically much faster than collisional rates. When these conditions are satisfied, external electromagnetic fields can be used to control

transport losses, enabling sustained thermonuclear reactivity in laboratory devices.

On the single-particle level, the motion of a charged particle with charge q and mass m in electric and magnetic fields \mathbf{E} and \mathbf{B} is governed by the Lorentz force

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

where \mathbf{v} is the particle velocity. In the simplest case with $\mathbf{E} = 0$ and uniform \mathbf{B} , particles perform fast gyromotion about magnetic field lines with cyclotron frequency $\Omega_c = |q|B/m$ and Larmor radius $\rho_L = v_\perp/\Omega_c$, while streaming freely along \mathbf{B} at the parallel velocity v_\parallel . The superposition of these motions produces helical trajectories that tend to keep particles close to magnetic field lines. This suggests that a strong magnetic field whose lines are closed in space could, in principle, confine a plasma.

However, a purely toroidal field \mathbf{B}_φ is not sufficient for this purpose. Field-strength gradients and toroidal curvature give rise to systematic guiding-centre drifts (such as curvature and ∇B drifts) that push ions and electrons in opposite vertical directions. The resulting charge separation drives an electric field that, together with the drifts, leads to net cross-field transport and, ultimately, to loss of confinement. To cancel these drifts and produce nested magnetic surfaces that effectively trap particles, a poloidal field B_θ must be superimposed so that the total field is helical. The pitch of the helical field lines is commonly characterised by the safety factor $q(r)$ which measures how many toroidal turns a field line makes per poloidal turn on a flux surface. Appropriate q -profiles are crucial for stability and transport: they help suppress magnetohydrodynamic (MHD) instabilities and reduce radial transport, thereby improving confinement.

The *tokamak* is one of the magnetic configurations that realises this concept. The toroidal field B_φ is generated by a set of toroidal-field (TF) coils encircling the vacuum vessel. The poloidal field B_θ is primarily produced by a toroidal current I_p driven in the plasma itself, which acts as the secondary of a transformer whose primary is a central solenoid. Additional poloidal-field (PF) coils provide vertical positioning and shaping control and help tailor the q -profile. The superposition of the TF and PF systems yields helical field lines that define a family of nested flux surfaces as shown in Fig. 1.3.

Tokamaks emerged in the late 1950s and 1960s from pioneering work in the Soviet Union by Tamm and Sakharov [7], followed by rapid inter-




Figure 1.3: Tokamak geometry and coordinate conventions. The plasma is organised in nested magnetic flux surfaces (coloured shells). A toroidal coordinate system (r, θ, φ) is used: r is the minor radius (distance from the magnetic axis), θ the poloidal angle around the small cross-section, and φ the toroidal angle around the machine. Together with the cylindrical frame (R, Z) , these define the major-radius and vertical directions. The zoom on the right shows a poloidal cross-section with closed flux contours (black), from the magnetic axis (red) to the last closed flux surface (LCFS).

national development. Decades of advances in magnet technology, plasma control, and plasma-facing components have made the tokamak the leading magnetic confinement concept and the basis of ITER, designed to explore burning-plasma physics and integrated operation at reactor scale, with a target fusion gain of $Q = 10$ and a plasma current of 15 MA [5].

1.4 Disruptions and runaway electrons

Paradoxically, the solution that makes a tokamak work, driving a large plasma current I_p to generate the poloidal field and helical flux surfaces, also constitutes one of its greatest hazards. The stored poloidal magnetic energy can cause an abrupt disruption and, possibly, the generation of runaway electrons.

1.4.1 Disruptions

Tokamak disruptions are off-normal events which cause the abrupt termination of the plasma in an uncontrolled way. Disruptions have multiple

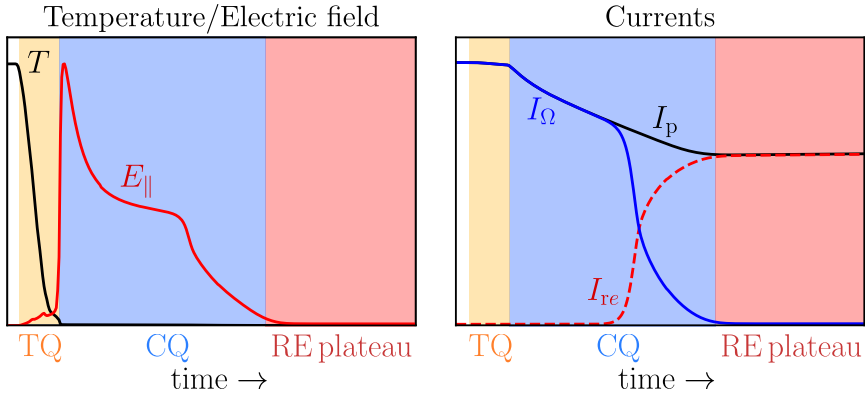


Figure 1.4: Illustration of the three main stages of a tokamak disruption. During the thermal quench (TQ) the electron temperature collapses, generating a large toroidal electric field E_{\parallel} that can accelerate electrons into the runaway regime. As the ohmic plasma current I_{Ω} decays during the current quench (CQ), a sufficiently large runaway population may form a runaway electron plateau, where the RE current I_{re} carries most of the total current I_p .

possible causes, but most arise when departures from stable operating conditions destabilise global magnetohydrodynamic (MHD) modes [8]. The dominant causes of disruptions are the excitation and nonlinear interplay of MHD instabilities, such as Neoclassical Tearing Modes (NTMs) and Resistive Wall Modes (RWMs), or parameter changes that push the equilibrium into an unstable regime, usually via modification of the current density profile as in the case of impurity accumulation [8].

Although the specific causes of a disruption can vary, the overall evolution of these events exhibits well-defined and recurrent features [4], as illustrated in Figure 1.4. The sequence begins with the *thermal quench* (TQ), triggered by the growth of large-scale MHD instabilities that lead to the ergodisation of the magnetic field lines. As a consequence, both particles and heat are rapidly transported to the first wall, resulting in a sudden and drastic drop in plasma temperature. The time scale of this event is usually $\lesssim 1$ ms.

The collapse of the electron temperature boosts electron–ion collisionality and raises the plasma resistivity since $\eta \propto T^{-3/2}$. The plasma thus becomes a poor conductor, yet a macroscopic toroidal current is still flowing. In the absence of external drive, the plasma–vessel system behaves like an inductive circuit and the current drops on a short resistive timescale. By flux conservation, the rapid fall of plasma current generates a large

toroidal electric field E_{\parallel} . In impurity-rich post-TQ conditions, this sequence leads to a rapid current drop on millisecond timescales.

The rapid transport of thermal energy to the first wall during the TQ imposes large transient heat loads on plasma-facing components, while halo currents exert significant electromagnetic forces on the vessel structure [8]. Additionally, the strong parallel electric field induced during the current quench can accelerate electrons to relativistic energies, creating a population of *runaway electrons* (REs). If the post-TQ runaway population, referred to as the *runaway seed* population, is large enough, it can be efficiently amplified through knock-on collisions during the CQ. The resulting runaway current can carry a large fraction of the total plasma current. When this occurs, the disruption enters a phase called *runaway plateau*.

On reactor-scale devices such as ITER, projections indicate that disruptions could generate runaway currents of several megaamperes. If not controlled, these highly energetic electron beams may impact the first wall or plasma-facing components, causing severe localised damage through melting, erosion and activation. To reduce these risks, future tokamaks rely on active disruption-mitigation schemes based on massive material injection (MMI), in which large quantities of mixed low-Z and high-Z impurities are injected into the plasma on a millisecond timescale. On current devices this is typically achieved through massive gas injection (MGI), where high-pressure gas is released from fast valves, while for ITER the baseline concept is shattered pellet injection (SPI), in which cryogenic impurity pellets are shattered into a spray of fragments just upstream of the plasma to improve penetration and assimilation. Understanding how such injections cool the plasma, modify the current quench and influence runaway generation is therefore a central element of ITER disruption mitigation. These aspects are investigated in this thesis through the numerical studies presented in Papers **II**, **III** and **IV**.

1.4.2 Runaway electrons

The electron runaway phenomenon can be understood by looking at the collisional friction force experienced by an electron moving in a plasma, as shown in Figure 1.5. This function is non-monotonic in the electron parallel velocity: it peaks around the thermal speed and then falls off like $1/v^2$. Once electrons are pushed beyond the velocity at which collisional drag can balance the accelerating electric force, they cross a threshold

in velocity space and are continuously accelerated, hence the term *run-away* [9]. Importantly, the electron acceleration is limited at very high energies due to synchrotron radiation, bremsstrahlung losses and partial screening effects [10–12].

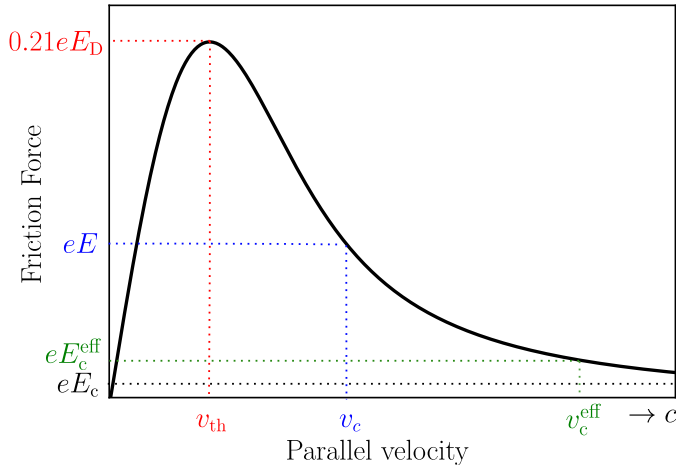


Figure 1.5: Collisional friction force versus parallel velocity v_{\parallel} experienced by electrons in a plasma. The drag peaks near the thermal speed v_{th} and decays as $\propto 1/v_{\parallel}^2$ at high v_{\parallel} . For a given electric field E , electrons above the critical velocity v_c experience a friction force smaller than the accelerating force eE and are therefore freely accelerated. In the presence of partially ionised impurities, bound electrons contribute additional friction, raising the effective critical field $E_c^{\text{eff}} > E_c$ and shifting the critical velocity to v_c^{eff} [12].

Two important electric fields can be defined when studying the runaway problem. The first one is referred to as the *Dreicer* field

$$E_D = \frac{e^3 n_e \ln \Lambda}{4\pi \epsilon_0^2 T_e}, \quad (1.4)$$

where $\ln \Lambda$ is the Coulomb logarithm, a parameter which quantifies the strength of Coulomb collisions. E_D is a characteristic electric field scale, defined by the condition that the accelerating force on a thermal electron equals the maximum collisional friction force. It sets the natural unit for electric fields in runaway physics, with all thermal electrons running away at $E \approx 0.21 E_D$ when the full friction force is accounted for [9].

The second one is the *critical* electric field

$$E_c = \frac{e^3 n_e \ln \Lambda}{4\pi \epsilon_0^2 m_e c^2}, \quad (1.5)$$

representing the relativistic limit for the existence of runaway electrons. The expression (1.5) was first derived by Connor and Hastie in 1975 [13]. This expression neglects radiation losses and incomplete screening in partially ionised plasmas. At high energy, electrons penetrate bound-electron shells and feel a larger effective ionic charge; together with bremsstrahlung and synchrotron, this raises the drag at high velocity. When these effects are included, the minimum electric field that can sustain runaway electrons becomes the *effective critical* electric field E_c^{eff} [12].

Runaway electrons are also a powerful source of electromagnetic radiation. Relativistic electrons spiralling in the tokamak magnetic field emit synchrotron and electron cyclotron radiation, and this emission can be exploited as a diagnostic of the underlying fast-electron distribution. Conventional electron cyclotron emission (ECE) systems are primarily designed to infer the bulk electron temperature under the assumption of a Maxwellian distribution. However, in the presence of a suprathermal tail the ECE spectrum becomes sensitive to the non-thermal population and, with an appropriate viewing geometry, may provide information on the nascent runaway seed. In this thesis, we make use of a dedicated Vertical ECE (VECE) system on TCV and of a synthetic ECE diagnostic to access the mildly relativistic part of the electron distribution and to characterise early runaway generation; these aspects are developed in **Paper I**.

1.5 Thesis outline

After this introduction, Chapter 2 develops the theoretical foundations of runaway electron generation, explaining how electrons can “run away”, and distinguishing between the primary and secondary generation mechanisms, along with runaway loss mechanisms.

Chapter 3 then turns to mitigation strategies aimed at preventing or suppressing runaway formation, with particular focus on shattered pellet injection (SPI), the leading technique for ITER, and on the processes governing the ablation and assimilation of injected material into the plasma. The numerical framework used to simulate disruptions and runaway generation is also introduced there.

Chapter 4 examines how runaway seed populations can be detected and characterised experimentally, presenting the development and modelling of the Vertical Electron Cyclotron Emission (VECE) diagnostic on the TCV tokamak, designed to measure cyclotron emission from suprathermal electrons. This is complemented by the modelling of suprathermal electron cyclotron emission through a synthetic diagnostic, combining wave propagation, radiation transport, and kinetic reconstruction of the electron distribution to enable direct comparison with experiments.

Chapter 5 concludes the thesis with a summary of the main findings and an outlook on future research directions, with particular relevance to forthcoming experiments and predictive modelling for ITER and next-generation fusion devices.

Chapter 2

Runaway electrons

The concept of “runaway” charged particles in an electric field is older than fusion research. Already in the 1920s, Wilson proposed that electrons in atmospheric electric fields could be continuously accelerated when the frictional force decreases sufficiently with energy, in order to explain high energy phenomena such as cosmic rays and lightning [14]. This qualitative idea of electrons “running away” was later placed on a firm kinetic footing by Dreicer, who solved the Fokker–Planck equation for a fully ionised plasma in a uniform electric field and derived both the critical “Dreicer field” and the corresponding runaway generation rate [15, 16]. Subsequent work by Gurevich and co-workers further elucidated the role of collisional diffusion and the formation of a diffusive flux across a critical momentum surface in phase space [17]. In the 1970s the theory was extended to relativistic energies and to parameters relevant for magnetically confined fusion plasmas. Connor and Hastie derived a relativistic expression for the critical field and showed how relativistic effects modify the runaway condition and generation rates in tokamak-like plasmas [13]. Around the same time, runaway electrons were being routinely observed in early tokamaks and other toroidal devices, both during current ramp-up and following plasma disruptions, and interpreted in terms of Dreicer-like acceleration in the large toroidal loop voltage [4, 18]. These observations established runaway electrons as a characteristic feature of tokamak plasmas.

2.1 The kinetic equation

Runaway electrons are inherently a phase-space phenomenon and, in principle, must be described by a kinetic equation for the electron distribution

function

$$f(\mathbf{x}, \mathbf{p}, t) \equiv \frac{dN}{d^3x d^3p}, \quad (2.1)$$

which gives the number of electrons in an infinitesimal phase-space volume $d^3x d^3p$ around position \mathbf{x} and momentum \mathbf{p} at time t .

The evolution of the distribution function for a species a is described by the kinetic equation

$$\frac{\partial f_a}{\partial t} + \frac{\mathbf{p}}{m_a \gamma} \cdot \frac{\partial f_a}{\partial \mathbf{x}} + q_a \left(\mathbf{E} + \frac{\mathbf{p}}{m_a \gamma} \times \mathbf{B} \right) \cdot \frac{\partial f_a}{\partial \mathbf{p}} = \sum_b C_{ab} \{f_a, f_b\}, \quad (2.2)$$

where \mathbf{E} and \mathbf{B} are the macroscopic electric and magnetic fields, $\gamma = \sqrt{1 + p^2}$ is the Lorentz factor, while C_{ab} represents the general Boltzmann collision operator which accounts for the effect of Coulomb interactions between species a and b . This operator is exact, but requires the solution of the full Boltzmann collision integral over the two-body cross section, which is outside the scope of the present work.

For the purposes of this thesis, the collisional dynamics will be described using the relativistic Fokker–Planck operator, which captures the cumulative effect of the many small-angle Coulomb collisions that dominate interactions between charged particles in a hot, fully ionised plasma. In particular, we model the interactions of species with both electrons and ions through the corresponding Fokker–Planck collision operator

$$C\{f\} = C_{ee}\{f\} + C_{ei}\{f\}, \quad (2.3)$$

which can be written in conservative form as

$$C\{f\} = -\nabla_{\mathbf{p}} \cdot [\mathbf{K}(\mathbf{p})f - \mathbb{D}(\mathbf{p}) \cdot \nabla_{\mathbf{p}} f], \quad (2.4)$$

where \mathbf{K} is the friction (or drag) vector and \mathbb{D} the momentum-space diffusion tensor. In a plasma near thermodynamic equilibrium these coefficients can be computed from the Landau collision operator and depend on the local density, temperature and effective charge of the plasma.

In the context of studying the runaway electron momentum space dynamics, it is convenient to adopt a spatially homogeneous description and treat the distribution function as a function of only \mathbf{p} and t , so that the Fokker-Planck equation can be rewritten as

$$\frac{\partial f}{\partial t} + \dot{\mathbf{p}} \cdot \nabla_{\mathbf{p}} f = C\{f\} + S\{f\}, \quad (2.5)$$

where $C\{f\}$ represents the Fokker-Planck collision operator and $S\{f\}$ contains additional sources and sinks. The quantity $\dot{\mathbf{p}} = d\mathbf{p}/dt$ contains all systematic forces acting on the electrons, such as acceleration by the inductive electric field and radiation reaction forces.

2.2 Runaway generation

In a kinetic description, runaway electrons are generally not a separate species but the high-energy, field-accelerated part of the electron distribution function. Their generation is therefore most naturally formulated as transport in momentum space. Neglecting spatial inhomogeneities, the kinetic equation can be written as a continuity equation in momentum space,

$$\frac{\partial f}{\partial t} + \nabla_{\mathbf{p}} \cdot \mathbf{J}_{\mathbf{p}} = S\{f\}, \quad (2.6)$$

where $\mathbf{J}_{\mathbf{p}}$ is the total momentum-space flux arising from systematic forces and Coulomb collisions. In a Fokker-Planck description

$$\mathbf{J}_{\mathbf{p}} = \dot{\mathbf{p}} f + \mathbf{K} f - \mathbb{D} \cdot \nabla_{\mathbf{p}} f, \quad (2.7)$$

where $\dot{\mathbf{p}}$ contains only the non-collisional forces acting on the electrons: acceleration by the inductive electric field and, when relevant, radiation-reaction forces, while \mathbf{K} and \mathbb{D} are the collisional friction vector and diffusion tensor introduced in (2.4).

Defining the runaway region Ω_{RE} as the part of momentum space above the instantaneous critical momentum p_c (and within a pitch-angle domain relevant for sustained acceleration), the runaway population

$$N_{\text{RE}}(t) = \int_{\Omega_{\text{RE}}} f d^3p \quad (2.8)$$

evolves according to

$$\frac{\partial N_{\text{RE}}}{\partial t} = - \oint_{\partial\Omega_{\text{RE}}} \mathbf{J}_{\mathbf{p}} \cdot \hat{\mathbf{n}} d\Sigma + \int_{\Omega_{\text{RE}}} S\{f\} d^3p, \quad (2.9)$$

so that runaway generation corresponds to a net flux of electrons from the thermal bulk into Ω_{RE} . Figure 2.1 illustrates the separation between the thermal bulk and the runaway region Ω_{RE} , where the separatrix $\partial\Omega_{\text{RE}}$ delineates the region where the electric field acceleration $\dot{\mathbf{p}}$ overcomes the collisional drag \mathbf{K} .

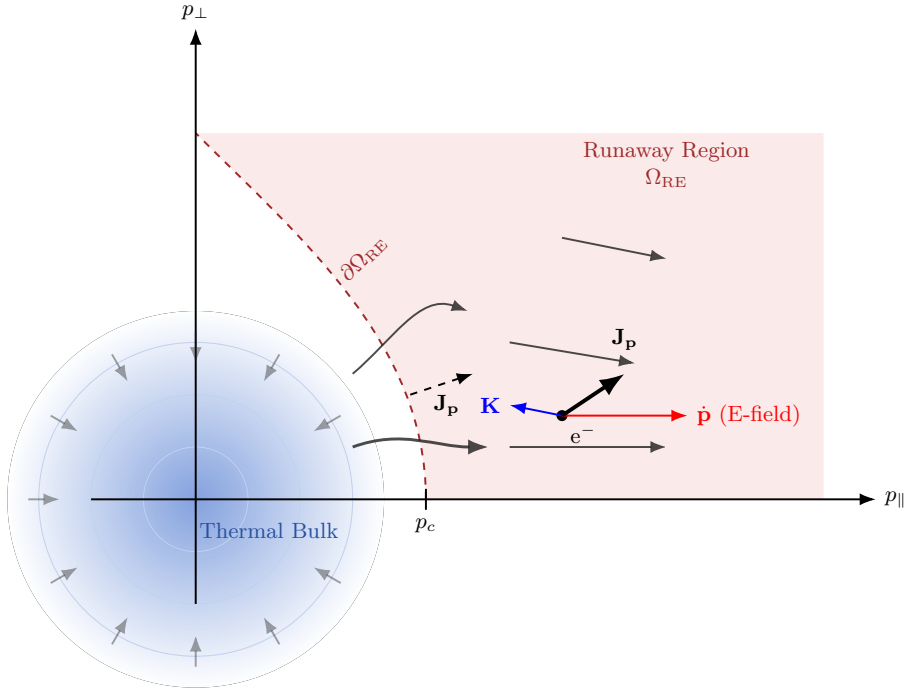


Figure 2.1: Schematic representation of runaway electron generation in momentum space $(p_{\parallel}, p_{\perp})$. The electron distribution is divided into a thermal bulk (blue) and a runaway region Ω_{RE} (red), separated by a critical momentum boundary $\partial\Omega_{\text{RE}}$ (dashed line). The total momentum-space flux \mathbf{J}_{p} is determined by the competition between the electric field acceleration $\dot{\mathbf{p}}$ and the collisional friction vector \mathbf{K} . Runaway generation corresponds to the net flux of electrons crossing the separatrix from the thermal bulk into Ω_{RE} .

In principle, the kinetic equation (2.6) for the electron distribution function, solved self-consistently with the macroscopic electric and magnetic fields and supplemented by a radial transport operator to describe spatial losses, captures all RE dynamics known to be important for practical applications. However, this description is itself an approximation: the radial transport operator is a phenomenological representation of effects, such as collisional or turbulent transport, that would arise naturally from more complete kinetic frameworks (e.g. drift-kinetic or gyrokinetic equations). While it provides the most complete description currently tractable, solving the kinetic equation is numerically demanding, partic-

ularly in disruption scenarios where rapid profile evolution, strong transport, and wide parameter scans are required. In such cases, reduced fluid models offer a computationally efficient alternative by replacing the full momentum-space flux across $\partial\Omega_{\text{RE}}$ with closed-form expressions for the primary sources and the avalanche growth rate, while retaining the same conceptual separation between seeding (flux into Ω_{RE}) and multiplication (growth within Ω_{RE}). The remainder of this chapter therefore describes runaway generation primarily in terms of such fluid models, where the runaway population is represented by a fluid density n_{RE} .

In the following, mechanisms that establish the initial RE population in Ω_{RE} (often called *seeds*) are referred to as *primary*. The source associated with large-angle knock-on collisions is denoted as the *secondary*, or *avalanche*, mechanism, since it exponentially amplifies an existing RE seed. All of these generation and loss mechanisms are implemented in DREAM (Disruption Runaway Electron Analysis Model) [19], the primary numerical framework used throughout this thesis.

2.2.1 Primary mechanisms

Primary generation mechanisms correspond to processes where electrons from the thermal bulk are directly promoted into the runaway region by the action of the electric field in combination with collisions, transient temperature changes or nuclear processes.

Dreicer generation. For a locally Maxwellian plasma, the classical Dreicer mechanism corresponds to a steady-state flux of electrons promoted from the thermal core across the runaway boundary by collisional diffusion in momentum space [15–17]. This boundary is set by the critical momentum p_c at which the accelerating electric force eE overcomes the collisional friction (Fig. 1.5), so that electrons diffusing past p_c are freely accelerated into the runaway region Ω_{RE} (Fig. 2.1).

In reduced (fluid) runaway models, Dreicer generation is represented by a *volumetric* source term for the runaway density,

$$\left(\frac{\partial n_{\text{RE}}}{\partial t}\right)_{\text{D}} \equiv \gamma_{\text{D}}(E_{\parallel}, T_e, n_e, Z_{\text{eff}}), \quad (2.10)$$

with units $\text{m}^{-3} \text{s}^{-1}$. In steady state, γ_{D} can be defined as the net momentum-

space flux into the runaway region,

$$\gamma_D = - \oint_{\partial\Omega_{RE}} \mathbf{J}_p \cdot \hat{\mathbf{n}} \, d\Sigma, \quad (2.11)$$

where $\hat{\mathbf{n}}$ is the outward normal of Ω_{RE} and \mathbf{J}_p is the total momentum-space flux.

Using a Fokker–Planck description with small-angle Coulomb collisions, Kruskal and Bernstein [20] and later Connor and Hastie [13] derived asymptotic expressions for the steady-state Dreicer generation rate. Written in the commonly used form,

$$\gamma_D = C \frac{n_e}{\tau_{ee}} \left(\frac{E_{\parallel}}{E_D} \right)^{-\frac{3}{16}(1+Z_{\text{eff}})h} \exp \left[-\lambda \frac{E_D}{4E_{\parallel}} - \sqrt{\eta(1+Z_{\text{eff}}) \frac{E_D}{E_{\parallel}}} \right], \quad (2.12)$$

where τ_{ee} is the electron–electron collision time, Z_{eff} is the effective ion charge, and C , h , λ and η are correction factors calculated by solving the relativistic Fokker–Planck equation in the limit of $E_{\parallel} \rightarrow E_c$ [13]. Equation (2.12) applies when the background plasma is stationary on the collisional timescale around the runaway threshold and the low-energy part of the distribution remains close to Maxwellian. It shows the strong (approximately exponential) sensitivity of Dreicer generation to E_{\parallel}/E_D : the source is negligible for $E_{\parallel} \ll E_D$, but can become significant already when E_{\parallel}/E_D reaches a few percent.

Hot–tail generation. A second primary mechanism which can provide a significant runaway seed in tokamak disruptions is *hot–tail* generation. It occurs when the bulk electron temperature collapses on a timescale comparable to, or shorter than, the collisional relaxation time of suprathermal electrons. For example, in a thermal quench, the electron temperature typically drops from a pre–disruption value $T_e \sim 5\text{--}10$ keV to a post–quench value $T_e \sim 10\text{--}100$ eV on a timescale $\tau_{TQ} \lesssim 1$ ms, while the electron density and effective charge can change substantially due to impurity ionisation. As a consequence of the strong velocity dependence of the collisional friction and scattering frequencies, the fast part of the pre–quench distribution cannot follow this rapid cooling and retains a non–Maxwellian “hot tail” which, in the presence of a rising inductive electric field, can feed the runaway region of phase space.

In a fluid description, hot–tail seeding can be represented in reduced form as a *boundary–flux* source into the runaway region. Denoting by $p_c(t)$

the instantaneous critical momentum defining the runaway boundary, the hot-tail source term for the runaway density is written [21]

$$\left(\frac{\partial n_{\text{RE}}}{\partial t}\right)_{\text{HT}} = -4\pi p_c^2 \dot{p}_c f_0(p_c, t), \quad (2.13)$$

where $f_0(p, t)$ is the angle-averaged hot-electron distribution during the cooling phase, evaluated at the boundary $p = p_c$. The factor $4\pi p_c^2$ arises from integrating the flux over the momentum-space surface at p_c , and \dot{p}_c accounts for the time dependence of the boundary as plasma parameters evolve. In the Lorentz (high- Z) limit, a pitch-angle averaged description yields the critical momentum p_c as the solution of

$$0 = \frac{1}{3} \left(\frac{E}{E_c}\right)^2 \frac{1}{1 + Z_{\text{eff}}} \frac{p_c^3}{\gamma} \frac{\partial f_0}{\partial p} + \frac{\gamma^2}{p^2} f_0. \quad (2.14)$$

This is the default reduced model implemented in the fluid-kinetic framework DREAM [19], which provides an analytical option for f_0 based on the hot-tail model of Smith and Verwichte [21]. In this approach, the survival of the tail during rapid cooling is captured by an isotropic slowing-down solution, which is inexpensive and therefore attractive for large scenario scans. However, it relies on the assumptions of strong pitch-angle scattering and negligible depletion of the suprathermal tail, and can therefore be sensitive to mechanisms that efficiently remove fast electrons during the thermal quench (e.g. stochastic magnetic transport), motivating benchmarking against higher-fidelity kinetic descriptions in regimes where hot-tail seeding is important.

Tritium β -decay. During nuclear operation, tritium nuclei present in the plasma undergo β^- -decay,



with a half-life $\tau_{\text{T}} \simeq 12.3$ yr, corresponding to a decay rate $\lambda_{\text{T}} = (\ln 2)/\tau_{\text{T}} \simeq 1.8 \times 10^{-9} \text{ s}^{-1}$. Each decay produces an electron with kinetic energy W drawn from a continuous β spectrum with endpoint $W_{\text{max}} \simeq 18.6$ keV and a broad maximum around a few keV [6]. For a given electric field and plasma composition, the critical Lorentz factor γ_c (or momentum p_c) defines the runaway boundary; a fraction of the emitted electrons are born with $\gamma > \gamma_c$ and can therefore enter the runaway region Ω_{RE} , essentially independently of the details of the bulk distribution. In more detailed

kinetic treatments, electrons born somewhat below the nominal threshold may also run away due to collisional diffusion in momentum space [22].

In the kinetic formulation of Eq. (2.5), tritium decay can be seen as an explicit source term $S_T\{f\}$ acting in momentum space. Assuming that the emission is isotropic in the tritium rest frame, S_T is nearly isotropic in pitch angle and localised at suprathermal energies $W \lesssim W_{\max}$. Its normalisation is fixed by

$$\int S_T\{f\} d^3p = \lambda_T n_T(t), \quad (2.16)$$

so that the total production rate of β electrons is $\lambda_T n_T$, where n_T is the tritium density. The detailed momentum dependence of S_T (in terms of the normalised β spectrum f_β) is not needed here and can be found in dedicated kinetic treatments [22].

Integrating the kinetic source over the runaway region Ω_{RE} yields the corresponding fluid runaway-generation rate

$$\left(\frac{\partial n_{\text{RE}}}{\partial t}\right)_T = \int_{\Omega_{\text{RE}}} S_T\{f\} d^3p \approx \lambda_T n_T F_\beta(\gamma_c), \quad (2.17)$$

where

$$F_\beta(\gamma_c) = \int_{\gamma_c}^{\gamma_{\max}} f_\beta(\gamma) d\gamma \quad (2.18)$$

is the fraction of β electrons born above the critical Lorentz factor γ_c associated with the runaway boundary [19, 23, 24]. The rate (2.17) is independent of the electric field except through the slowly varying threshold $\gamma_c(E_{\parallel})$.

Compton scattering. Fusion neutrons in a reactor-scale device activate the first wall and surrounding structures, producing a substantial flux of γ -rays toward the plasma [25]. Energetic γ -photons crossing the plasma can transfer part of their energy to background electrons through Compton scattering; if the scattered electron energy exceeds the runaway threshold, a new runaway is created. This provides a volumetric, field-independent seed that depends mainly on the γ -ray flux and on the total number of target electrons, both free and bound [23, 25].

The kinematics of a single Compton event between a photon of initial energy E_γ and a stationary electron is described by the standard Compton formula. For a scattering angle θ between the incoming and outgoing

photon directions, the final photon energy is

$$E'_\gamma = \frac{E_\gamma}{1 + \frac{E_\gamma}{m_e c^2} (1 - \cos \theta)}, \quad (2.19)$$

and the electron kinetic energy gain is

$$W_e = E_\gamma - E'_\gamma, \quad \gamma = 1 + \frac{W_e}{m_e c^2}, \quad p = \sqrt{\gamma^2 - 1}. \quad (2.20)$$

Only scattering events with $\gamma \geq \gamma_c$ (or $p \geq p_c$) contribute to runaway generation. Their probability is determined by the Klein–Nishina differential cross section $d\sigma/d\Omega(E_\gamma, \theta)$ and by the range of scattering angles that place the scattered electron above the runaway threshold [23, 25].

In the kinetic formulation of Eq. (2.5), Compton scattering enters as an explicit source term $S_\gamma\{f\}$. For an approximately isotropic photon field with spectral flux density $\Gamma(E_\gamma)$ (number of photons per unit area, time and energy), and denoting by n_{tot} the total electron density (including bound electrons), its normalisation is fixed by

$$\int S_\gamma\{f\} d^3p = n_{\text{tot}} \int \Gamma(E_\gamma) \sigma(E_\gamma) dE_\gamma, \quad (2.21)$$

where $\sigma(E_\gamma)$ is the total Klein–Nishina cross section. The detailed momentum dependence of S_γ (in terms of the differential cross section and the photon spectrum) is obtained numerically in dedicated studies [22, 23].

As for tritium decay, a fluid description of Compton seeding is obtained by integrating the kinetic source over the runaway region,

$$\left(\frac{\partial n_{\text{RE}}}{\partial t} \right)_\gamma = \int_{\Omega_{\text{RE}}} S_\gamma\{f\} d^3p \approx n_{\text{tot}} \int \Gamma(E_\gamma) \sigma_{\text{run}}(E_\gamma, p_c) dE_\gamma, \quad (2.22)$$

where $\sigma_{\text{run}}(E_\gamma, p_c)$ is the Klein–Nishina cross section integrated over scattering angles that produce electrons with $p \geq p_c$ [19, 23]. This “runaway-effective” cross section depends on plasma parameters only through the critical momentum $p_c(E_\gamma, n_e, Z_{\text{eff}})$ that defines the runaway boundary.

2.2.2 Secondary mechanism – avalanche generation

In addition to primary mechanisms, runaway electrons can be produced by *secondary*, or *avalanche*, generation: a relativistic runaway electron undergoes a *large-angle* Coulomb collision (Møller scattering) with a thermal electron and transfers sufficient momentum that both of the outgoing

electrons cross the runaway boundary. This “knock-on” mechanism was identified early on by Sokolov and later by Jayakumar *et al.* as a potentially dominant multiplication channel in high-current devices [26, 27]. Since each runaway electron can generate further runaways, the process is intrinsically multiplicative and leads to an exponential increase of the runaway density once a seed population is present [28].

In fluid models, avalanche generation is represented by a source term proportional to the existing runaway density,

$$\left(\frac{\partial n_{\text{RE}}}{\partial t}\right)_{\text{ava}} = \Gamma_{\text{ava}}(E_{\parallel}, n_{\text{e}}, Z_{\text{eff}}, \dots) n_{\text{RE}}, \quad (2.23)$$

where Γ_{ava} is the avalanche growth rate. In the fully ionised, relativistic limit, Rosenbluth and Putvinski obtained the expression

$$\Gamma_{\text{ava}}^{(\text{RP})} \simeq \frac{1}{\tau_{\text{c}} \ln \Lambda_{\text{c}}} \sqrt{\frac{\pi}{3(Z_{\text{eff}} + 5)}} \left(\frac{E_{\parallel}}{E_{\text{c}}} - 1\right), \quad \frac{E_{\parallel}}{E_{\text{c}}} \gg \sqrt{1 + Z_{\text{eff}}}, \quad (2.24)$$

where τ_{c} is the relativistic electron collision time, $\ln \Lambda_{\text{c}}$ the relativistic Coulomb logarithm, and E_{c} the Connor–Hastie critical field [13, 28]. This expression captures the approximately linear scaling $\Gamma_{\text{ava}} \propto E_{\parallel}$ at large fields and the vanishing of avalanche multiplication as $E_{\parallel} \rightarrow E_{\text{c}}$.

In plasmas containing partially ionised impurities, large-angle collisions on *bound* electrons and partial screening effects must be included, which modify both the effective critical field and the collision frequencies experienced by fast electrons [29, 30]. In that case, the avalanche growth rate can be written in the form [31]

$$\Gamma_{\text{ava}} = \frac{e}{m_{\text{e}} c \ln \Lambda_{\text{c}}} \frac{n_{\text{e,tot}}}{n_{\text{th}}} \frac{E_{\parallel} - E_{\text{c}}^{\text{eff}}}{\sqrt{4 + \bar{\nu}_{\text{s}}(p_{\text{c}}^*) \bar{\nu}_{\text{D}}(p_{\text{c}}^*)}}, \quad (2.25)$$

where $n_{\text{e,tot}}$ is the total density of target electrons (free plus bound), n_{th} the free thermal electron density, $E_{\text{c}}^{\text{eff}}$ an effective critical field including screening effects, and $\bar{\nu}_{\text{s}}$ and $\bar{\nu}_{\text{D}}$ are the normalised slowing-down and deflection frequencies evaluated at an effective critical momentum p_{c}^* . The prefactor $n_{\text{e,tot}}/n_{\text{th}}$ accounts for the increased number of available targets when partially ionised impurities are present, while the dependence on $E_{\text{c}}^{\text{eff}}$ [12] and the collision frequencies accounts for the enhanced drag and pitch-angle scattering. In disruption–mitigation scenarios with substantial amounts of injected high- Z material, the increased number of targets often dominates, leading to a higher avalanche multiplication rate than in the fully ionised case [31].

2.3 Runaway electron losses

The final population of runaway electrons is determined by the competition between the generation mechanisms described in Section 2.2 and various loss channels. In the context of tokamak disruptions, confinement is degraded significantly and losses occur primarily through two distinct mechanisms: radial transport across magnetic flux surfaces driven by magnetic perturbations, and direct scrape-off against plasma-facing components due to the macroscopic vertical motion of the plasma column.

2.3.1 Transport losses

During the thermal quench, the rapid cooling of the plasma is typically precipitated by the destabilisation of magnetohydrodynamic (MHD) modes. These instabilities break the toroidal symmetry of the magnetic field, leading to magnetic islands and, in many cases, a stochastic field-line topology. In such a field, magnetic field lines wander radially as they traverse the torus; since charged particles stream primarily along field lines, this induces strong radial transport of heat and particles [32].

A common reduced description models the radial dynamics as diffusion with a coefficient D_{rr} . In the collisionless limit relevant for relativistic electrons, the Rechester–Rosenbluth [32] estimate gives

$$D_{\text{rr}} \approx \pi q R_0 |v_{\parallel}| \left(\frac{\delta B}{B} \right)^2, \quad (2.26)$$

where q is the safety factor, R_0 the major radius, v_{\parallel} is the parallel electron velocity, and $\delta B/B$ is the relative amplitude of the magnetic perturbations.

However, energetic runaway electrons do not remain tied to a single perturbed field line. Because of curvature and ∇B drifts, their guiding-centre orbit deviates from the original field line over a parallel correlation length. When this drift-induced displacement becomes comparable to or larger than the correlation length of the magnetic perturbation, the electron samples effectively decorrelated fluctuations and the radial transport is reduced at high momentum [33]. A simple momentum-dependent diffusion coefficient capturing this decorrelation, used in reduced disruption modelling, is [34, 35]

$$D_{\text{RE}}(p) = D_{\text{rr}} \frac{p}{1+p^2} \xrightarrow{p \gg 1} D_{\text{rr}} \frac{1}{p}, \quad (2.27)$$

written in terms of the normalised momentum $p \equiv |\mathbf{p}|/(m_e c)$. This form leaves the transport of the thermal bulk to a separate heat-transport treatment, while capturing the reduced stochastic transport of suprathermal/runaway electrons during the thermal quench, which is critical for determining whether an initial seed survives to be amplified during the current quench.

In a fluid description of the runaway density n_{RE} , transport losses are described by the divergence of the radial runaway flux Γ_{RE} ,

$$\left(\frac{\partial n_{\text{RE}}}{\partial t}\right)_{\text{transport}} = -\frac{1}{V'} \frac{\partial}{\partial r} (V' \Gamma_{\text{RE}}), \quad (2.28)$$

where V' is the radial Jacobian. Consistent with the phase-space definition $f = dN/(d^3x d^3p)$, the runaway particle flux is obtained by integrating the local diffusive flux over pitch angle ξ and momenta above a runaway threshold p_* ,

$$\Gamma_{\text{RE}} = - \int_{p_*}^{\infty} \int_{-1}^1 2\pi p^2 D_{\text{RE}}(p) \frac{\partial f_{\text{RE}}}{\partial r} d\xi dp. \quad (2.29)$$

In disruption scenarios where the runaway population is dominated by avalanche multiplication, $f_{\text{RE}}(p, \xi)$ may be approximated by an analytical avalanche distribution, which enables an efficient closure of Eq. (2.29) without fully resolving the runaway electron distribution function [34].

2.3.2 Scrape-off losses

In disruptions of highly elongated plasmas such as ITER, an additional and often decisive loss mechanism is the scrape-off of runaway electrons caused by the vertical plasma motion. As the plasma current I_p decays during the current quench, inductive coupling to passive conducting structures can destabilise the vertical position of elongated equilibria and drive a vertical displacement event (VDE) [36]. Unlike stochastic transport, which can be modelled as diffusion across flux surfaces, a VDE causes macroscopic flux surfaces to intersect material boundaries. Runaway electrons residing on flux surfaces that become open are then lost essentially on the parallel transit time to the wall,

$$\tau_{\text{loss}} \sim \frac{R}{c} \sim 10^{-8} - 10^{-7} \text{ s}, \quad (2.30)$$

which is effectively instantaneous compared to the millisecond-to-100 ms disruption evolution timescales.

Modelling this process self-consistently generally requires 2D MHD simulations, at least. To enable efficient integrated simulations in DREAM, a reduced scrape-off model was developed in **Paper III** [37]. The model exploits that, in reactor-scale devices such as ITER, the resistive time of the conducting wall is typically long compared to the current-quench duration. As a consequence, the poloidal magnetic flux at the last closed flux surface (LCFS) remains approximately constant on the disruption timescale, and the position of the LCFS can be estimated by solving [37, 38]

$$\psi_p(r_{\text{LCFS}}(t), t) \approx \psi_p(a, t = 0), \quad (2.31)$$

where a is the initial minor radius and r is the (fixed) flux-surface label used in the 1D transport model. The scrape-off losses are implemented as a localised sink acting only outside the LCFS,

$$\left(\frac{\partial n_{\text{RE}}}{\partial t} \right)_{\text{scrape-off}} = -\frac{n_{\text{RE}}}{\tau_{\text{loss}}} \Theta(r - r_{\text{LCFS}}(t)), \quad (2.32)$$

so that runaway electrons on open flux surfaces are removed on the fast timescale τ_{loss} . Since τ_{loss} is orders of magnitude shorter than all other relevant timescales, it may be chosen on the order of the numerical timestep (e.g. $\tau_{\text{loss}} \sim 10\Delta t$) without loss of accuracy and without imposing excessive time resolution requirements.

This reduced model is primarily intended for scenarios where scrape-off overlaps with the runaway generation phase, for which the constant- ψ_p approximation is valid; it is generally not suitable for describing the termination of long-lived runaway plateaus on times comparable to the wall resistive time [37]. The implications of scrape-off for ITER disruption scenarios are investigated in detail in **Papers II** and **III**.

Chapter 3

Disruption mitigation

Disruptions in reactor-scale tokamaks pose three main challenges: (i) the rapid release of thermal and magnetic energy, driven by the ergodisation of flux surfaces illustrated in Fig. 3.1, which can produce intolerable heat loads on plasma-facing components; (ii) strong electromagnetic forces arising from halo currents and fast current decay; and (iii) the production of RE that can evolve into a confined beam and, upon termination, deposit energy in a highly localised region. In ITER, as well as in other advanced devices such as STEP and SPARC, these hazards justify the use of an active Disruption Mitigation System (DMS) based on *massive material injection* (MMI), realised either through Massive Gas Injection (MGI) or Shattered Pellet Injection (SPI). The purpose of MMI is not to *prevent* a disruption after it has been initiated, but rather to *control* the disruption dynamics so that energy is dissipated mainly through radiation, the current quench proceeds on acceptable timescales, and the post-quench electric field remains sufficiently low to limit significant runaway multiplication.

This chapter is structured as follows. Section 3.1 presents the two principal forms of massive material injection (MGI and SPI), together with the injection-scheme variants relevant to ITER. Then, Section 3.2 develops, step by step, the integrated disruption model applied in **Papers II–IV**. Finally, Section 3.3 introduces the DREAM framework, which is one of the main numerical tools used in this thesis.




Figure 3.1: Evolution of the Poincaré plots for a positive triangularity plasma for varying magnitude of magnetic perturbations $\delta B/B$. At $\delta B/B = 0$ all the flux surfaces are intact and closed; at $\delta B/B = 5 \times 10^{-4}$ the magnetic field starts to form islands with some local ergodicity; at high magnetic perturbation $\delta B/B = 5 \times 10^{-3}$, the magnetic confinement is mostly lost and the magnetic field lines are mostly ergodic.

3.1 Massive material injection

The primary class of actuators used for disruption mitigation is *massive material injection* (MMI), meaning the rapid introduction of a large particle inventory (typically a combination of low- Z fuel and radiating impurities) into the plasma just before or during a disruption [39]. Two forms of MMI are employed in current experiments and considered in the design of future devices: *massive gas injection* (MGI) and *shattered pellet injection* (SPI).

3.1.1 Massive gas injection

In MGI, a pre-filled high-pressure reservoir is connected to the vacuum vessel through a fast valve, releasing a short, intense gas pulse into the edge plasma [39]. MGI is mechanically straightforward and capable of delivering large particle inventories on millisecond timescales. Its main drawback for disruption mitigation in hot, large-scale plasmas is *assimilation*: the injected gas quickly ionises and becomes magnetised near

the plasma edge, which can decrease deep penetration and produce pronounced spatial asymmetries in both deposition and radiation [39, 40]. These limitations become more important as plasma temperature and size increase, providing the main motivation for pellet-based methods under ITER-scale conditions.

3.1.2 Shattered pellet injection

SPI is a specialised form of conventional cryogenic pellet injection, developed specifically to transport MMI inventories further into the plasma [41]. In this method, a cryogenic pellet is produced in a cooled barrel and then pneumatically accelerated. Immediately before entering the vacuum vessel, the pellet is deliberately fragmented, typically through impact within a bent “shatter” tube or against a dedicated shattering element, thereby generating a *spray* of a number of shards [41]. This fragmentation increases the effective surface area of the injected material, which enhances ablation and reduces the likelihood that a large intact pellet will traverse the plasma and deposit material in an overly localised manner. By construction, the resulting shard-size distribution, and therefore the penetration depth and deposition profile, can be influenced through the choice of pellet material composition, injection velocity, and shatter geometry [41].

This injection scheme was first established experimentally for disruption mitigation on DIII-D [42] and has subsequently advanced from a proof-of-principle method to the baseline concept for the ITER DMS [43, 44]. Compared to MGI, SPI provides a more effective route towards high impurity assimilation and more spatially distributed radiation in reactor-relevant plasmas, which is why it plays a central role in the modelling and scenario optimisation examined in this thesis.

Although SPI denotes the fragmentation-based delivery method itself, the resulting mitigation outcome depends strongly on *what* is injected (the material composition), *how* it is injected (the staging and timing), and *where* it is injected (the number and spatial distribution of injection locations). The principal scheme variants considered in present devices and in ITER-oriented studies are:

(i) **Single-stage injection.** In the single-stage approach, one shattered pellet contains both low- Z fuel, intended to increase the plasma density and collisionality, and a high- Z radiating impurity, intended to convert thermal and magnetic energy into radiation and thereby reduce conducted

heat fluxes. For ITER, the baseline concept focuses primarily on mixtures based on hydrogen isotopes with neon, and in some studies argon, as the radiating species [43–45]. The main practical advantage of this scheme is its operational simplicity.

(ii) Staggered injection. An alternative strategy is to separate the functions of dilution and radiation by injecting material in two successive steps. In this case, a predominantly low- Z pellet is injected first to raise the density and shorten the characteristic timescales of collisional processes, after which an impurity-rich pellet is injected following a controlled delay to maximise radiative dissipation once the plasma has already been diluted. The purpose of this approach is to influence both the thermal-quench dynamics and the conditions for RE generation by shaping the temperature and density evolution during mitigation [46]. This staged concept, often referred to as staggered injections, can be of particular interest for reactor-scale devices, where RE avoidance and mitigation are key drivers in DMS design.

(iii) Multi-pellet single-stage injection. More generally, beyond two distinct stages, an SPI system may in principle inject a *train* of pellets with prescribed temporal spacing in order to control the effective deposition rate and reduce sensitivity to timing jitter relative to the disruption onset. Such schemes represent a natural extension in systems with multiple injectors that can be triggered independently [44, 45].

SPI facilities are now present on several tokamaks as part of the experimental program supporting validation of the ITER DMS. The first demonstrations of SPI for disruption mitigation were performed on DIII-D [42]. Since then, SPI systems and SPI-based disruption mitigation experiments have also been implemented on JET [47], ASDEX Upgrade [48], and KSTAR [49], among others, making possible systematic investigations of penetration, assimilation, radiation peaking, and RE mitigation across a range of machine sizes and plasma operating regimes.

Looking ahead, SPI is expected to remain a key element of disruption mitigation in next-step and reactor-oriented tokamaks. In fact, the ITER DMS is explicitly founded on SPI, using multiple toroidally distributed injectors to improve radiation symmetry and to provide flexibility in both the injected material mixtures and the injection timing [43–45]. Beyond ITER, disruption and RE mitigation requirements are also being assessed

for future power-plant concepts, including STEP [50] and the Chinese reactor development pathway toward CFETR [51].

3.2 Disruption modelling

Predictive assessment of disruption mitigation by massive material injection requires a *coupled* description of (i) how material enters the plasma, (ii) how it is ablated and ionised, (iii) where the resulting particles and radiative power are deposited, and (iv) how the evolving composition, temperature and electric field feed back into runaway-electron generation. In integrated disruption simulations, the material injection is modelled as a set of discrete “shards” whose trajectories determine localised sources of particles and energy. These sources are then mapped to flux-surface averaged quantities through a deposition kernel, and evolved self-consistently with the plasma. The present section describes each component of this coupled model in turn.

3.2.1 Material injection

In SPI, a cryogenic pellet is fragmented in a shatter tube before entering the vessel. From a modelling perspective, the main output of the shattering process is an ensemble of shards, characterised by radii $\{r_{p,k}\}_{k=1}^{N_s}$, velocities $\{\mathbf{v}_{p,k}\}$ initial propagation directions that define a plume opening angle and composition. Because fragmentation is inherently stochastic, it is natural to describe the shattering process by a probabilistic model for the shard-size distribution and, where appropriate, for correlated shard velocities, and then to carry out sensitivity studies over multiple realizations.

A useful statistical description of fragment sizes is obtained by assuming that the pellet breaks through random fragmentations, which yields a shard-radius probability density of the form

$$P(r_p) = k_p^2 r_p K_0(k_p r_p), \quad (3.1)$$

where K_0 denotes the modified Bessel function of the second kind. The characteristic inverse length k_p is determined by requiring consistency with the total injected inventory,

$$k_p = \left(\frac{N_{\text{inj}}}{6\pi^2 n_p N_s} \right)^{-1/3}, \quad (3.2)$$

where N_{inj} is the total number of injected atoms, n_p is the number density of the solid pellet material, and N_s is the number of shards. Fragmentation models of this type have been employed in SPI studies of disruption mitigation to generate shard ensembles for subsequent ablation and deposition modelling, and they can be combined with machine-specific plume models describing the velocity spread and opening angle [52–54].

Once a shard ensemble has been specified, the simplest treatment of the shard kinematics is to model each shard as a ballistic projectile until complete ablation:

$$\mathbf{x}_{p,k}(t) = \mathbf{x}_{p,k}(t_0) + \mathbf{v}_{p,k}(t_0)(t - t_0), \quad (3.3)$$

where t_0 is the time at which the shard ensemble is initialized, typically near the plasma boundary. The distribution of $\mathbf{v}_{p,k}$ is then chosen to represent a conical plume with a prescribed effective opening angle and speed dispersion [53, 54].

This model neglects acceleration arising from asymmetric ablation, often referred to as the “pellet rocket effect” [55]. Nevertheless, it is sufficient to capture the dominant parameters governing penetration and assimilation, namely the characteristic shard size, the injection speed, and the shard composition.

3.2.2 Material ablation and deposition

After entering the hot plasma, each shard rapidly forms a dense and cold neutral cloud. Plasma electrons incident on the shard are cooled and scattered within this cloud, which *shields* the solid surface and thereby reduces the ablation rate well below what would be predicted from a direct free-streaming estimate. This effect constitutes the physical foundation of the *neutral gas shielding* (NGS) model [56, 57] and of its more recent extensions developed for mixed-material pellets and disruption-relevant conditions [58].

In the modelling framework used in this thesis, it is convenient to describe the ablation process through an evolution equation for the shard radius $r_{p,k}$ in the form of a radius recession law. One practical expression, used in reduced kinetic SPI modelling, is

$$\dot{r}_{p,k} = -\frac{\lambda(X_k)}{4\pi r_{p,k}^2 \rho_{\text{dens}}} \left(\frac{q_{\text{in},k}}{q_{\text{ch}}}\right)^{1/3} \left(\frac{E_{\text{in},k}}{E_{\text{ch}}}\right)^{7/6} \left(\frac{r_{p,k}}{r_{\text{ch}}}\right)^{4/3}, \quad (3.4)$$

where ρ_{dens} is the solid density, X_k denotes the shard composition, for example the neon fraction in a D–Ne pellet, and $\lambda(X_k)$ is a dimensionless prefactor that accounts for the dependence of the ablation physics on composition.

The normalisation parameters are a characteristic shard radius $r_{\text{ch}} = 2$ mm, a characteristic electron energy flux $q_{\text{ch}} = n_{\text{ch}} \sqrt{2T_{\text{ch}}/(\pi m_e)}$, and a characteristic electron energy $E_{\text{ch}} = 2T_{\text{ch}}$, with the reference temperature and density set to $T_{\text{ch}} = 2000$ eV and $n_{\text{ch}} = 10^{20}$ m⁻³, respectively [59]. For a D–Ne pellet with deuterium fraction $X = N_{\text{D}_2}/(N_{\text{D}_2} + N_{\text{Ne}})$, the composition-dependent prefactor takes the form [35]

$$\lambda(X) = 2.99 \times 10^{23} [27.08 + \tan(1.49 X)] \text{ s}^{-1}. \quad (3.5)$$

The quantities $q_{\text{in},k}$ and $E_{\text{in},k}$ denote, respectively, the incident electron energy flux and a characteristic incident electron energy at the shard position. For an approximately isotropic electron distribution $f_e(\mathbf{p})$, they are defined as

$$q_{\text{in}} = \frac{1}{4} \int d^3p m_e c^2 (\gamma - 1) v f_e(\mathbf{p}), \quad (3.6)$$

$$E_{\text{in}} = \frac{2}{n_{\text{free}}} \int d^3p m_e c^2 (\gamma - 1) f_e(\mathbf{p}), \quad n_{\text{free}} = \int d^3p f_e(\mathbf{p}), \quad (3.7)$$

where m_e is the electron mass, v the speed corresponding to momentum \mathbf{p} , γ the relativistic factor, and n_{free} the free-electron density.

For a Maxwellian plasma, q_{in} reduces, up to a numerical factor, to the one-sided thermal flux $n_e v_T T_e/4$, while $E_{\text{in}} = 2T_e$. The more general definitions in Eqs. (3.6)–(3.7) additionally retain the contribution from suprathermal electrons, which may become important if a hot tail persists during the thermal quench.

The ablation rate expressed in atoms per unit time follows directly from the rate of pellet-volume loss,

$$\dot{N}_k = -4\pi r_{p,k}^2 n_p \dot{r}_{p,k}, \quad (3.8)$$

where n_p is the solid atomic number density. For a pellet composed of several atomic species s with atomic fractions $f_{s,k}$, for example D and Ne, the source term for each species is then $\dot{N}_{s,k} = f_{s,k} \dot{N}_k$.

Cross-field drift. Ablated material is ionised and initially transported along magnetic field lines. However, experiments have long indicated that

pellet ablation clouds can undergo a systematic drift toward the low-field side. This *plasmoid drift* is commonly interpreted as an interchange-like response of a dense, high- β filament, or “plasmoid”, embedded in a curved magnetic field: grad- B and curvature effects generate a net diamagnetic current that must be balanced by polarisation and parallel closure currents, thereby producing a poloidal electric field and an associated radial $E \times B$ drift [60–62].

A reduced current-closure model captures the essential physics by imposing current continuity,

$$\nabla \cdot (\mathbf{j}_{\nabla B} + \mathbf{j}_{\text{pol}} + \mathbf{j}_{\parallel}) = 0, \quad (3.9)$$

where a standard estimate for the grad- B /curvature-driven current within the plasmoid is

$$\mathbf{j}_{\nabla B} \simeq -\frac{2p}{BR_m} \hat{\mathbf{y}}, \quad (3.10)$$

with p denoting the plasmoid pressure, B the magnetic-field strength, R_m an effective magnetic-curvature radius, and $\hat{\mathbf{y}}$ a direction that is approximately poloidal in the local cross-section. The polarisation current can be written as

$$\mathbf{j}_{\text{pol}} = \frac{\rho}{B^2} \frac{\partial \mathbf{E}_{\perp}}{\partial t}, \quad (3.11)$$

with ρ the mass density. The resulting poloidal electric field $E_y(t)$ then gives rise to a radial $E \times B$ drift velocity,

$$v_x(t) = \frac{E_y(t)}{B}, \quad \Delta x(t) = \int_0^t \frac{E_y(t')}{B} dt'. \quad (3.12)$$

In the long-time limit, $\Delta x(\infty)$ represents the net radial displacement of the deposited material relative to a picture in which the material simply follows the magnetic field lines. Inclusion of plasmoid drift corresponds to taking

$$r_k(t) = r_{p,k}(t) + \Delta x_k(t), \quad (3.13)$$

with $\Delta x_k(t)$ obtained from the drift model in Eq. (3.12). This formulation isolates the dominant *systematic* outward bias in assimilation and enables its impact on mitigation performance to be quantified. A semi-analytical plasmoid-drift model of this kind was formulated and benchmarked for pellet-ablation conditions in [63]. In the present thesis, the same model plays a central role since it was implemented in DREAM in **Paper IV**, validated against SPI disruptions in ASDEX Upgrade, and extrapolated to ITER-like conditions [64, 65].

3.2.3 Density and charge-state evolution

The ablation model of the previous section gives us localised particle sources, while the evolution of the charge-state distribution of all ion species self-consistently with ionisation, recombination, and transport determines the plasma composition. This affects the energy balance through radiation, the evolution of the electric field through conductivity, and the evolution of the runaway electrons through their charge and number of bound electron targets.

Once the material is ablated, the dense plasmoid material drifts across the magnetic field lines and expands along them on a timescale $\lesssim 1$ ms, comparable to that of thermal quench but much shorter than that of current quench. In the 1D framework used here, homogenization is assumed to happen instantaneously, i.e., the material is homogenized in the poloidal direction instantly, with ablated material entering as a flux surface average, a similar approximation to that used in other disruption mitigation studies [66, 67].

If n_{ij} is the density of charge state i of ion species j , then the evolution of the charge-state distribution is governed by

$$\begin{aligned} \frac{\partial n_{ij}}{\partial t} = & \left. \frac{\partial n_{ij}}{\partial t} \right|_{\text{ioniz/rec}} + S_{ij}(r, t) \\ & + \frac{1}{V'} \frac{\partial}{\partial r} \left[V' \left(-A_{ij} n_{ij} + D_{ij} \frac{\partial n_{ij}}{\partial r} \right) \right], \end{aligned} \quad (3.14)$$

where the three terms on the right-hand side represent, respectively, ionisation and recombination, the ablation source, and advective–diffusive cross-field transport prescribed to approximate the enhanced transport observed in stochastic magnetic fields.

The ionisation and recombination term is computed by integrating the time-dependent rate equations,

$$\begin{aligned} \left. \frac{\partial n_{ij}}{\partial t} \right|_{\text{ioniz/rec}} = & I_{i-1,j} n_{i-1,j} n_M - I_{ij} n_{ij} n_M \\ & + R_{i+1,j} n_{i+1,j} n_M - R_{ij} n_{ij} n_M, \end{aligned} \quad (3.15)$$

where $I_{ij}(T_M, n_M)$ and $R_{ij}(T_M, n_M)$ are the ionisation and recombination rate coefficients, interpolated from tabulated values in the ADAS [68] and AMJUEL databases, and n_M, T_M denote the thermal (Maxwellian) electron density and temperature. The free-electron density n_{free} (including

both Maxwellian and non-Maxwellian populations) is determined at each time step by the quasi-neutrality constraint

$$n_{\text{free}} = \sum_{ij} i n_{ij}. \quad (3.16)$$

Importantly, the freshly ablated material is deposited into the neutral charge state of the relevant species. When the plasmoid drift model is active (as in **Papers II** and **IV**), the ablated material has already been ionised to at least the singly charged state by the time the drift terminates; it is then deposited in the singly charged state [63].

3.2.4 Energy balance and thermal quench

Within the integrated 1D disruption-modelling framework adopted in this thesis, the thermal quench is not modelled as a fully resolved MHD event. Instead, it is treated as a rapid transition phase in which transport and radiation deplete the plasma thermal energy, driving the plasma into a cold and highly resistive state. In this framework, the evolution of the bulk “cold” Maxwellian electron population is described through a flux-surface-averaged energy-balance equation. This population is characterised by density n_M and temperature T_M , with thermal energy density $W_M = \frac{3}{2}n_M T_M$. For the studies presented in **Papers II–IV**, this energy balance is solved self-consistently as

$$\begin{aligned} \frac{\partial W_M}{\partial t} = & \frac{3}{2V'} \frac{\partial}{\partial r} \left(n_M V' D_W \frac{\partial T_M}{\partial r} \right) \\ & - \left. \frac{\partial W_M}{\partial t} \right|_{\text{line}} - \left. \frac{\partial W_M}{\partial t} \right|_{\text{ioniz}} + \left. \frac{\partial W_M}{\partial t} \right|_{\text{ohm}} \\ & + \left. \frac{\partial W_M}{\partial t} \right|_{\text{col}} - \left. \frac{\partial W_M}{\partial t} \right|_{\text{brem}} - \left. \frac{\partial W_M}{\partial t} \right|_{\text{halo}}, \end{aligned} \quad (3.17)$$

where the right-hand-side terms account for radial heat transport caused by stochastic magnetic fields, line-radiation losses, ionisation losses, ohmic heating, collisional energy exchange with ions, bremsstrahlung losses, and, when relevant, parallel heat losses in the halo region during vertical displacement events. The first six contributions are included in all disruption simulations. The final term, by contrast, is only activated when the plasma experiences a vertical displacement such that the flux surfaces become open. A brief account of these terms is given below, while a full

description of their physical meaning and numerical implementation can be found in [19].

During a disruption, MHD activity can stochasticise the magnetic field, so that electrons follow perturbed field lines and their radial motion may be approximated as diffusive. For a single test electron moving along the field with parallel velocity v_{\parallel} , the Rechester–Rosenbluth model gives the radial diffusion coefficient

$$D_{\text{RR}}(v) = \pi q R_0 v_{\parallel} (\delta B/B)^2,$$

where q is the safety factor, R_0 the major radius, and $\delta B/B$ the relative magnetic perturbation amplitude [32]. The corresponding *heat* diffusion coefficient entering the energy-balance equation is obtained by taking the heat-flux moment of D_{RR} over a Maxwellian distribution. In other words, the velocity-space average is weighted by the energy transported by each particle relative to the mean, which introduces the factor $v^2/v_T^2 - 3/2$. This gives

$$D_W = \frac{1}{\pi^{3/2} v_T^3 T_M} \int \frac{m_e v^2}{2} \left(\frac{v^2}{v_T^2} - \frac{3}{2} \right) D_{\text{RR}}(v) \exp\left(-\frac{v^2}{v_T^2}\right) dv, \quad (3.18)$$

where $v_T = \sqrt{2T_M/m_e}$ is the electron thermal speed.

Because a self-consistent calculation of $\delta B/B$ would require a fully three-dimensional MHD description, its time evolution is prescribed in the 1D model employed in this thesis.

In this modelling framework, stochastic transport is not assumed to be active throughout the disruption, and a criterion is therefore required to determine the onset of the thermal quench following material injection. One possible choice is to activate it when the local temperature falls below a critical value, $T_{\text{crit}} \sim 10$ eV, which may indicate the development of sufficiently sharp ohmic current density gradients to destabilise resistive MHD modes. An alternative is to use the passage of impurity-doped pellet shards across the $q = 2$ surface as the trigger. The consequences of adopting these different onset criteria are examined in **Paper II**. Once activated, $\delta B/B$ amplitudes in the range 0.1–1% are typically sufficient to reduce the temperature to about 100 eV within a few milliseconds, consistent with expected ITER thermal-quench timescales [25].

At temperatures below roughly 100 eV, the cooling is primarily driven

by radiation and ionisation losses,

$$\left. \frac{\partial W_M}{\partial t} \right|_{\text{line}} = n_M \sum_{ij} n_{ij} L_{ij}(T_M, n_M), \quad (3.19)$$

$$\left. \frac{\partial W_M}{\partial t} \right|_{\text{ioniz}} = n_M \sum_{ij} n_{ij} E_{ij}^{\text{ioniz}} I_{ij}(T_M, n_M), \quad (3.20)$$

where L_{ij} denotes the line-radiation rate coefficients, taken from the ADAS database, I_{ij} the ionisation rate coefficients appearing also in Eq. (3.15), and E_{ij}^{ioniz} the corresponding ionisation energies from the NIST database. As the plasma cools, these loss channels are eventually offset by ohmic heating,

$$\left. \frac{\partial W_M}{\partial t} \right|_{\text{ohm}} = \sigma_{\parallel} E_{\parallel}^2, \quad (3.21)$$

where σ_{\parallel} is the Spitzer parallel conductivity and E_{\parallel} is the representative flux-surface-averaged parallel electric field, defined in Section 3.2.5. The competition between radiative cooling and ohmic heating determines the post-thermal-quench equilibrium temperature, which typically lies in the range of a few eV.

Energy exchange between electrons and ions is described by the collisional term

$$\left. \frac{\partial W_M}{\partial t} \right|_{\text{col}} = \frac{3 n_M m_e}{m_i \tau_e} (T_i - T_M), \quad (3.22)$$

where m_i is the ion mass, T_i the ion temperature, and τ_e the electron collision time. Although this contribution is usually subdominant during the thermal quench, it is important for maintaining energy conservation between the two species.

Lastly, an additional loss channel can appear during plasma vertical displacements. In this case, the plasma column moves vertically as the elongated equilibrium loses vertical stability while the plasma current decays. Once the displacement becomes sufficiently large, the outermost flux surfaces intersect the vessel wall and become open, thereby forming a *halo* region. Electrons in this region stream along open field lines and transport energy directly to the wall, creating a parallel heat-loss channel that may dominate the conducted heat load in disruptions with a low radiated power fraction.

Following the model introduced in Ref. [69], this effect is represented by the sink term

$$\left. \frac{\partial W_M}{\partial t} \right|_{\text{halo}} = \frac{4}{3} \gamma_{\text{sh}} \frac{n_e c_s T_e}{L_{\parallel}} \Theta(r - r_{\text{LCFS}}), \quad (3.23)$$

where Θ is the Heaviside step function, r_{LCFS} is the instantaneous position of the last closed flux surface, and $L_{\parallel} = 2\pi q R_0$ is the parallel connection length of the flux tube. The sheath heat-transmission coefficient γ_{sh} measures the energy carried through the sheath per escaping particle. In the simulations, the value $\gamma_{\text{sh}} = 8$ is adopted, consistent with standard sheath theory for a hydrogen plasma with $T_e \approx T_i$. The ion sound speed is taken as

$$c_s = \sqrt{\frac{T_e + \gamma T_i}{m_i}},$$

with $\gamma = 5/3$ corresponding to an adiabatic transformation with isotropic pressure, and m_i taken as the mass of the majority ion species. This loss term is only active in the open-field-line region, $r > r_{\text{LCFS}}$, but in low-radiation disruptions it can generate substantial heat loads on the vessel wall, consistent with observations from JET with the ITER-like wall.

3.2.5 Electric field and currents

The rapid cooling described in the previous section leads to a sharp increase in the plasma resistivity, since $\eta \propto T^{-3/2}$. As a result, a large toroidal electric field is induced as the system responds to approximately conserve the poloidal magnetic flux. Within the 1D framework, the parallel electric field is therefore obtained self-consistently from the evolution of the poloidal flux $\psi_p(r, t)$.

Expressed in terms of the flux-surface label r and the toroidal flux ψ_t , the mean-field induction equation takes the form

$$\frac{\partial \psi_p}{\partial t} = V_{\text{loop}} + \mu_0 \frac{\partial}{\partial \psi_t} \left(\psi_t \Lambda \frac{\partial j_{\parallel}}{\partial \psi_t B} \right), \quad (3.24)$$

where $V_{\text{loop}} = 2\pi \langle E \cdot B \rangle / \langle B \cdot \nabla \phi \rangle$ is the loop voltage. The final term, proportional to Λ , is included to model the flattening of the current-density profile caused by magnetic stochasticity [70, 71]. In the limit $\Lambda = 0$, Eq. (3.24) reduces to the usual Faraday induction law. The poloidal flux

is related to the total parallel current density j_{\parallel} through Ampère's law,

$$2\pi\mu_0\langle B \cdot \nabla\phi \rangle \frac{j_{\parallel}}{B} = \frac{1}{V'} \frac{\partial}{\partial r} \left(V' \left\langle \frac{|\nabla r|^2}{R^2} \right\rangle \frac{\partial\psi_p}{\partial r} \right), \quad (3.25)$$

which closes the relation between the current profile and the evolving magnetic flux.

At the plasma boundary, the poloidal flux is coupled to the surrounding conducting structures through the mutual inductance M_{ew} . The wall flux then evolves according to

$$\psi_{\text{wall}} = -L_{\text{ext}}(I_p + I_{\text{wall}}), \quad \frac{\partial\psi_{\text{wall}}}{\partial t} = R_{\text{wall}} I_{\text{wall}}, \quad (3.26)$$

where $L_{\text{ext}} = \mu_0 R_0 \ln(R_0/b)$ is the external inductance, $R_{\text{wall}} = L_{\text{ext}}/t_{\text{wall}}$ is the wall resistance, and t_{wall} is the resistive wall time.

The total parallel current density is written as the sum of an ohmic and a runaway contribution,

$$j_{\parallel} = j_{\text{ohm}} + j_{\text{RE}}, \quad (3.27)$$

where the ohmic current is carried by the thermal Maxwellian population,

$$\frac{j_{\text{ohm}}}{B} = \sigma_{\parallel} \frac{\langle E \cdot B \rangle}{\langle B^2 \rangle}, \quad (3.28)$$

while the runaway current is given by

$$\frac{j_{\text{RE}}}{B} = ec \frac{n_{\text{RE}}}{\langle B \rangle}, \quad (3.29)$$

assuming that the runaway electrons stream along the magnetic field at approximately the speed of light. The evolution of the runaway density n_{RE} is discussed in the next section.

3.2.6 Runaway electrons

In DREAM, the time evolution of the runaway-electron density n_{RE} during a mitigated disruption is obtained by combining the generation mechanisms introduced in Chapter 2 with the relevant loss channels, all evaluated using the self-consistent electric field and plasma profiles described

in the previous sections:

$$\begin{aligned} \frac{\partial n_{\text{RE}}}{\partial t} = & \left. \frac{\partial n_{\text{RE}}}{\partial t} \right|_{\text{D}} + \left. \frac{\partial n_{\text{RE}}}{\partial t} \right|_{\text{HT}} + \left. \frac{\partial n_{\text{RE}}}{\partial t} \right|_{\text{T}} \\ & + \left. \frac{\partial n_{\text{RE}}}{\partial t} \right|_{\gamma} + \left. \frac{\partial n_{\text{RE}}}{\partial t} \right|_{\text{ava}} + \left. \frac{\partial n_{\text{RE}}}{\partial t} \right|_{\text{transp}} \\ & + \left. \frac{\partial n_{\text{RE}}}{\partial t} \right|_{\text{scrape-off}}, \end{aligned} \quad (3.30)$$

where the terms on the right-hand side represent, respectively, Dreicer generation, hot-tail generation, tritium β -decay, Compton scattering, avalanche multiplication, radial transport losses, and scrape-off losses. The corresponding expressions were introduced in Sections 2.2.1–2.3 of Chapter 2.

The coupling between Eq. (3.30) and the electric-field evolution (3.24) is a key element of the disruption dynamics. As n_{RE} increases, the runaway current (3.29) replaces an increasing fraction of the decaying ohmic current. This slows the overall current decay and, in turn, reduces the induced electric field. This negative feedback ultimately sets the final runaway plateau current in a self-consistent simulation.

The fluid model also requires a prescription for the critical momentum p_c that separates the thermal bulk from the runaway region in momentum space. In a fully ionised plasma subject to a uniform electric field, p_c follows from the balance between electric-field acceleration and collisional drag. In realistic post-disruption plasmas, however, this balance is modified by both pitch-angle scattering and partial screening. The expression adopted in DREAM is obtained by matching asymptotic solutions of the bounce-averaged kinetic equation in the limits of strong pitch-angle scattering and near-threshold electric field [19]:

$$p_c^* = m_e c \left(\frac{[\bar{\nu}_s(p_c^*) \bar{\nu}_D(p_c^*) + 4\bar{\nu}_s^2(p_c^*)] E_c^2}{f_p (\tilde{E}_{\parallel} - E_c^{\text{eff}})^2} \right)^{1/4}, \quad (3.31)$$

where $\bar{\nu}_s$ and $\bar{\nu}_D$ are the normalised slowing-down and deflection frequencies, describing respectively the rate of parallel-momentum loss and the rate of pitch-angle scattering. These quantities are evaluated at an effective critical momentum p_c^* , which accounts self-consistently for partial screening. The quantity E_c^{eff} is the effective critical electric field [12], while f_p is the passing-particle fraction, reflecting the fact that only passing electrons can undergo sustained acceleration by the parallel electric field.

3.3 Numerical tools

The disruption simulations presented in **Papers II, III and IV** of this thesis were carried out using the DREAM code [19]. DREAM is a fluid–kinetic modelling framework that solves the coupled system of equations described in the preceding sections, including SPI particle sources, charge-state and density evolution, self-consistent energy balance, electric-field diffusion, and runaway-electron dynamics. Starting from prescribed pre-disruption profiles and injection parameters, it evolves the plasma self-consistently through the disruption.

In the context of this thesis, the DREAM fluid model is employed. Here, the electron distribution is divided into a Maxwellian bulk, characterised by n_M and T_M , and a runaway population, characterised by n_{RE} and assumed to travel at the speed of light. The runaway density is then evolved using the analytic generation and loss terms collected in Eq. (3.30). This mode is computationally efficient, typically requiring $\lesssim 1$ CPU-hour per simulation, and is therefore well suited for the large parameter scans performed in **Paper II**.

The full equation system is discretised using a finite-volume method on a one-dimensional radial grid in configuration space and, in kinetic mode, on the corresponding (p, ξ_0) grid in momentum space. Time integration is performed with a backward-Euler scheme. At each time step, the resulting implicit nonlinear system is solved by Newton iteration using analytically constructed Jacobians [19]. The finite-volume formulation ensures exact conservation of particle number and energy, up to machine precision, in the absence of source terms and edge losses.

Chapter 4

Suprathermal electron cyclotron emission

The detection of runaway electrons is notoriously difficult in their initial phase as “seeds”, particularly when their energy is only mildly relativistic with a phase-space density that is a small perturbation on the overall Maxwellian distribution. Nevertheless, this phase is critical to the future behaviour of runaway electrons, as once a seed population is established and conditions are appropriate for avalanche-type multiplication [28], the growth in runaway current can be significant and difficult to control.

The majority of conventional methods for detecting runaways are intrinsically biased toward ultra-relativistic electrons with energy in the range of MeV to tens of MeV, either because this is when a given radiation mechanism is significant or because background signals are more significant at lower energy. Hard X-ray cameras with bremsstrahlung detection rely on runaways with energy in the MeV range to produce copious numbers of high-energy photons [72–74]. Synchrotron-based diagnostics can provide powerful constraints on the behaviour of the high-momentum, high-pitch part of the distribution function [75–77], but are again intrinsically biased toward ultra-relativistic electrons. Diagnostics involving wall interactions, such as loss monitors, infrared thermography, power balance, etc., are invaluable tools for understanding loss channels [39] but provide little direct constraint on the confined seed population.

By contrast, the mildly relativistic domain of interest for runaway seeding generally corresponds to an electron energy in the range $10\text{--}10^2\text{keV}$, in which the radiation from bremsstrahlung processes tends to be less significant and may be spectrally contaminated by background radiation, and in

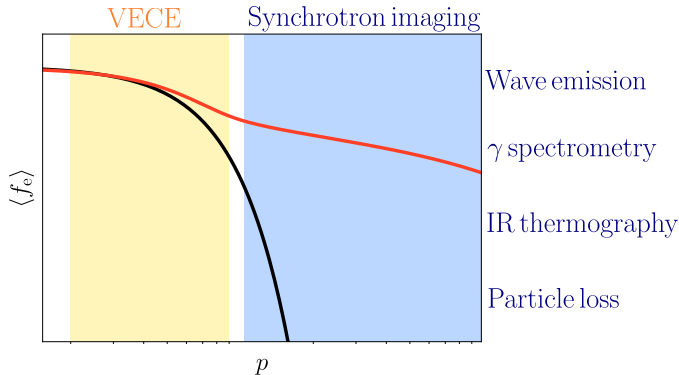


Figure 4.1: Schematic coverage of fast-electron diagnostics as a function of electron momentum p , expressed as qualitative sensitivity to different regions of the distribution function $f_e(p)$. At low to intermediate momenta, VECE constrains the mildly relativistic, confined electron population in the seed and early suprathermal regime, where the emission is optically thick and the spectrum encodes the shape of the distribution function. At intermediate to high momenta, HXR/ γ spectrometry provides energy-resolved information on the suprathermal tail. At high momenta, synchrotron imaging and wave emission diagnostics (e.g. whistlers) constrain the ultra-relativistic population, while particle loss detectors and IR thermography characterise losses and power deposition at material surfaces.

which neutron detection does not offer an independent handle on the seed. To overcome this blind spot in the standard set of diagnostics, illustrated schematically in Fig. 4.1, specialised diagnostics have been developed that project the phase space information of suprathermal electrons into some form of spectral measurement.

The Vertical Electron Cyclotron Emission (VECE) diagnostic realises this principle by looking at the plasma from above, i.e., along a nearly vertical line of sight, and measuring the spectrum of ECE radiation over a broad frequency band, in order to project the measured radiation frequency into the information of the electron energy rather than the emission radius. By doing so, the VECE diagnostic fills an important gap in the existing set of diagnostics, as it allows for the measurement of suprathermal electrons in the confined, mildly relativistic domain of interest, relevant for the seeding and the initial formation of the suprathermal tail. The principle was pioneered in the 1980s and 1990s on the PLT [78], Alcator C [79], Tore Supra [80, 81], and DIII-D [82] devices, and has recently been pursued on TCV [83, 84].

4.1 Vertical Electron Cyclotron Emission

Diagnosing the mildly relativistic stage of fast-electron formation using ECE requires a viewing geometry that lets frequency encode electron energy rather than position. In a conventional horizontal LOS, the strong radial variation of the magnetic field couples ECE frequency to both radius and electron energy, so a spectral change cannot be unambiguously attributed to the suprathermal tail rather than a shift in the emitting layer [85, 86], obscuring the suprathermal tail. VECE adopts a *vertical* LOS so that, at nearly fixed major radius, the ECE frequency is governed mainly by the harmonic number n and the relativistic factor γ . In the weakly relativistic limit, the cyclotron frequency scales as $f_{\text{ECE}} \propto n/\gamma$; when a single harmonic dominates, the spectrum can therefore be interpreted as an energy spectrum, with frequency broadening mapping primarily to energy broadening through $\mathcal{E} = m_e c^2 (\gamma - 1)$.

This energy selectivity originates from the relativistic cyclotron–Doppler resonance condition

$$\omega - k_{\parallel} v_{\parallel} = \frac{n\Omega_{ce}}{\gamma}, \quad (4.1)$$

where $\Omega_{ce} = eB/m_e$ is the non-relativistic electron cyclotron frequency, k_{\parallel} is the component of the wavevector along the magnetic field, and v_{\parallel} is

the electron parallel velocity. For a near-vertical view, k_{\parallel} is small and the observed frequency shift is set chiefly by γ and the pitch-angle distribution, rather than by radial field variation. Deviations from strict energy encoding arise from the finite beam width, refraction-induced changes in k_{\parallel} , and overlap between neighbouring harmonics at sufficiently large energies.

The reference diagnostic for this thesis is the TCV VECE [83, 84], whose poloidal viewing geometry is shown in Fig. 4.2. It measures




Figure 4.2: Poloidal view of the VECE diagnostic lines of sight in TCV [83].

in the 78–114 GHz band with $10\ \mu\text{s}$ time resolution and $\sim 750\ \text{MHz}$ channel bandwidth. A wire-grid polariser splits X- and O-modes, the optics form a Gaussian beam with a $\sim 3\ \text{cm}$ waist at the midplane, and the LOS terminates on a highly absorbing dump at the vessel bottom [87] to suppress wall reflections. In this configuration VECE is primarily sensitive to down-shifted X3–X4 emission from fast electrons; within the observing band the accessible energies extend from tens to a few hundred keV (up to $\sim 250\ \text{keV}$ before harmonic overlap becomes important) [83].

Operation requires attention to constraints that define the effective “clean window” of the measurement. Refraction can bend the LOS so that it misses the dump, reintroducing thermal background through multipath reflections; density and viewing frequency must therefore be chosen so that the ray bundle remains well controlled [83, 84]. Channels that intersect strongly opaque thermal layers (notably thermal X2 in many configurations) become contaminated and are excluded. Conversely, B_0 can be tuned so that usable channels avoid thermal resonances while retaining sensitivity to the desired suprathermal energies. At the high-energy end, overlap between adjacent harmonics limits unique energy assignment; at the low-energy end, the thermal brightness sets the dynamic range for detecting a sparse non-thermal tail.

4.2 Suprathermal ECE modelling

Interpreting VECE measurements requires a forward model that connects a *phase-space* electron distribution function to the *spectral* signal recorded by the radiometer. Reconstructing a 3D distribution from a limited number of frequency channels is an ill-conditioned inverse problem, while the forward problem, given f_e and equilibrium, to predict the detected intensity, is well posed and can be validated. This forward-modelling philosophy, adopted in Paper I [88] and earlier works on JET [89] and ASDEX Upgrade [90], is the approach followed throughout this thesis.

The modelling strategy constructs a synthetic VECE diagnostic that combines: (i) a kinetic calculation of the suprathermal electron distribution function, (ii) a geometrical-optics description of propagation along the diagnostic LOS, and (iii) a radiation transport solver that includes emission and (re)absorption along the ray. This section summarises the physical ingredients and approximations relevant to the mildly relativistic regime targeted by VECE.

4.2.1 Wave propagation

When the vacuum wavelength is short compared with equilibrium scale lengths, ECE modelling can be cast in the geometrical-optics (WKB) approximation [91, 92]. Wave propagation is then described by rays whose trajectories are governed by a local dispersion relation

$$\mathcal{D}(\omega, \mathbf{k}, \mathbf{x}) = 0, \quad (4.2)$$

where \mathcal{D} is determined by the cold-plasma dielectric tensor and by the selected polarisation branch. The corresponding ray equations, written in terms of the determinant Λ of the dispersion tensor and used by the C3PO code [93], take the Hamiltonian form

$$\frac{d\mathbf{r}}{dt} = - \left. \frac{\partial \Lambda / \partial \mathbf{k}}{\partial \Lambda / \partial \omega} \right|_{\Lambda=0}, \quad \frac{d\mathbf{k}}{dt} = \left. \frac{\partial \Lambda / \partial \mathbf{r}}{\partial \Lambda / \partial \omega} \right|_{\Lambda=0}. \quad (4.3)$$

Ray tracing of this kind has been extensively applied to lower-hybrid and electron-cyclotron wave propagation [94–97].

In the present context, ray tracing serves two purposes. First, density gradients refract the rays, altering k_{\parallel} and thereby shifting the Doppler-broadened resonance condition (4.1), while also displacing the spatial region that contributes to the measured signal. Refraction can further cause the ray bundle to miss the viewing dump, reintroducing reflected thermal radiation into the line of sight. A synthetic diagnostic therefore must use ray tracing to determine validity windows (frequency band, density regime, field ramp) within which the intended near-vertical propagation and isolation from wall reflections are guaranteed [84, 88].

Second, the ray trajectory supplies the local quantities that enter the emission and absorption coefficients: the magnetic field B (and hence the cyclotron frequency Ω_{ce}), the refractive-index vector $\mathbf{N} = c\mathbf{k}/\omega$ with its parallel and perpendicular components, and the branch-dependent polarisation factors. Together, these determine which regions of momentum space contribute to the emission at each observed frequency.

4.2.2 Radiation transport

Along a ray trajectory $\mathbf{x}(s)$ parameterised by path length s , the specific spectral intensity I_{ω} obeys the radiative transfer equation [91, 92]

$$\frac{d}{ds} \left[\frac{I_{\omega}(s)}{N_{\omega,r}^2(s)} \right] = \frac{1}{N_{\omega,r}^2(s)} [j_{\omega}(s) - \alpha_{\omega}(s) I_{\omega}(s)], \quad (4.4)$$

where j_ω is the emissivity, α_ω the absorption coefficient for the selected mode and harmonic, and $N_{\omega,r}$ the ray refractive index. When only the intensity at the antenna is required, the prefactor $N_{\omega,r}^2$ can be set to unity — a simplification shown to hold also for non-thermal distributions [92]. The transfer equation then reduces to $dI_\omega/ds = j_\omega - \alpha_\omega I_\omega$, whose formal solution between a source point s_0 and the receiver at s_1 reads

$$I_\omega(s_1) = I_\omega(s_0) e^{-\tau_\omega(s_0 \rightarrow s_1)} + \int_{s_0}^{s_1} j_\omega(s) e^{-\tau_\omega(s \rightarrow s_1)} ds, \quad (4.5)$$

with optical depth $\tau_\omega(s \rightarrow s_1) \equiv \int_s^{s_1} \alpha_\omega(s') ds'$. If the line of sight terminates on a viewing dump of negligible reflectivity, as is typically the case for the TCV VECE diagnostic, the boundary term vanishes ($I_\omega(s_0) = 0$).

For suprathermal ECE the emission and absorption coefficients must be evaluated from the full electron distribution function, which may be decomposed as

$$f_e = f_M + \tilde{f}, \quad (4.6)$$

where f_M is the bulk Maxwellian (with spatially varying n_e, T_e) and \tilde{f} the suprathermal component. In the spectral windows exploited by VECE, the two populations play distinct roles: the thermal bulk can dominate absorption and hence set the optical thickness, while the suprathermal tail dominates emission at frequencies down-shifted from the cold resonance, where the Maxwellian contribution is exponentially small. This separation is precisely what makes VECE viable: the plasma must remain sufficiently transparent for suprathermal emission to reach the antenna, yet optically thick enough for the signal to be interpretable through the resonance condition (4.1).

4.2.3 Emission and absorption coefficients

The radiative transfer equation (4.4) requires explicit expressions for j_ω and α_ω . Following standard practice in synthetic ECE modelling, momentum components (p_\parallel, p_\perp) are normalised to $m_e c$, and the cyclotron resonance condition (4.1) takes the compact form

$$\gamma - p_\parallel N_\parallel - \frac{n}{\bar{\omega}} = 0, \quad \bar{\omega} \equiv \frac{\omega}{\Omega_{ce}}, \quad (4.7)$$

with $N_\parallel \equiv ck_\parallel/\omega$ evaluated along the ray.

The harmonic- n contributions to the absorption coefficient and the emissivity are obtained by integrating over momentum space [98, 99]:

$$\alpha_\omega^{(n)}(s) = -\frac{2\pi^2}{c\omega} \omega_{p,0}^2 \iint |\Theta|^2 \hat{R}_n f_e \delta\left(\gamma - p_\parallel N_\parallel - \frac{n}{\bar{\omega}}\right) \frac{p_\perp}{\gamma} dp_\perp dp_\parallel, \quad (4.8)$$

$$j_\omega^{(n)}(s) = \frac{m_e c^2}{4\pi c^3} N_{\omega,r}^2 \omega_{p,0}^2 \omega \iint |\Theta|^2 f_e \delta\left(\gamma - p_\parallel N_\parallel - \frac{n}{\bar{\omega}}\right) \frac{p_\perp}{\gamma} dp_\perp dp_\parallel, \quad (4.9)$$

where $\omega_{p,0}$ is the electron plasma frequency at the local density, Θ a mode-dependent polarisation factor incorporating finite-Larmor-radius effects, and $N_{\omega,r}$ the ray refractive index for the relevant branch. The differential operator

$$\hat{R}_n \equiv \frac{n\Omega_{ce}}{\omega p_\perp} \frac{\partial}{\partial p_\perp} + N_\parallel \frac{\partial}{\partial p_\parallel} \quad (4.10)$$

acts on f_e in the absorption integral but is absent from the emissivity. This asymmetry reflects a basic physical distinction: emission samples f_e directly on the resonance surface, whereas absorption depends on velocity-space gradients of f_e through \hat{R}_n . For a Maxwellian the two are linked by detailed balance, recovering the Kirchhoff relation; for a non-thermal distribution, however, the gradient structure of \tilde{f} can make the absorption spectrum qualitatively different from the emission spectrum.

The total coefficients are obtained by summing over harmonics,

$$j_\omega = \sum_{n \in \mathcal{N}} j_\omega^{(n)}, \quad \alpha_\omega = \sum_{n \in \mathcal{N}} \alpha_\omega^{(n)}, \quad (4.11)$$

where the set \mathcal{N} is chosen to cover the VECE frequency band, typically the third and fourth X-mode harmonics, which dominate the down-shifted emission in the regimes of interest.

A compact way to connect the distribution function to the measured signal is through a Green's function $D_\omega(s, p_\parallel)$ [88], defined such that

$$I_\omega = \iint f_e[r(s), p_\parallel] D_\omega(s, p_\parallel) ds dp_\parallel, \quad (4.12)$$

where $r(s)$ is the radial location as a function of the ray path coordinate s , while the kernel D_ω encapsulates the ray geometry, the resonance condition, the polarisation factors, and the optical-depth weighting from

Eqs. (4.5) – (4.9). The integrand $f_e D_\omega$, referred to as the *birthplace distribution*, identifies which regions in physical and momentum space dominate the detected signal at a given frequency, providing a direct map from features in f_e to signatures in the VECE spectrum, as demonstrated in Paper I.

4.3 Numerical tools

The suprathermal ECE modelling described above is implemented through an integrated toolchain whose components are summarised in this section. The guiding principle is to separate kinetic modelling of f_e , wave propagation, and radiative transfer into distinct codes and then couple them consistently through a synthetic diagnostic. Figure 4.3 illustrates the data flow.




Figure 4.3: Forward-modelling pipeline for synthetic VECE emission, coupling the LUKE Fokker–Planck solver, the C3PO ray-tracing code, and the YODA synthetic diagnostic.

Figure 4.3 illustrates the forward-modelling pipeline. LUKE solves the Fokker–Planck equation for the electron distribution function f_e given equilibrium inputs ($T_e, n_e, V_{\text{loop}}$) and, optionally, the EC wave deposition profile provided by C3PO. C3PO traces rays along the VECE line of sight, supplying the local magnetic field B and refractive index N_{\parallel} to YODA, which evaluates the emission, absorption, and radiative transfer to produce the synthetic intensity I_ω . The three numerical tools and their coupling are briefly described in the following sections.

4.3.1 LUKE: Fokker–Planck solver

The electron distribution function is computed with LUKE [100], a bounce-averaged, relativistic Fokker–Planck code that evolves f_e under four operators: Coulomb collisions \mathcal{C} , parallel electric-field acceleration \mathcal{E} , quasilinear EC diffusion \mathcal{Q} (evaluated with the ray-tracing code C3PO [93]), and ra-

dial transport \mathcal{S} . The formulation is numerically conservative, ensuring particle-number preservation on each flux surface [100].

In the electron-cyclotron current-drive (ECCD) scenarios analysed in [84, 88], the non-inductive current modifies the total current profile and thereby the induced loop voltage, which in turn governs the ohmic acceleration shaping the suprathermal tail. To capture this feedback, LUKE evolves the flux-surface-averaged toroidal induction equation, coupling the radial redistribution of loop voltage to changes in the current density in a self-consistent way, essential for reproducing VECE signals [88].

4.3.2 C3PO: ray tracing

Wave propagation between the plasma and the antenna is modelled with C3PO [93], which integrates the cold-plasma dispersion relation for the selected O- or X-mode branch following Eq. (4.3). The VECE antenna pattern is represented as a Gaussian beam with waist $w_0 \simeq 3$ cm at the tokamak midplane [83], discretised into a bundle of rays whose number is chosen so that the calculated ECE intensity is converged. Diffraction is neglected within the geometrical-optics approximation; refraction is retained through the spatially varying dielectric response. Rather than launching rays from every emitting point toward the antenna, the model exploits electromagnetic reciprocity: rays are traced *from* the antenna into the plasma and onward to the viewing dump, and the resulting beam pattern weights the emission collected at the antenna [88, 93].

4.3.3 YODA: synthetic ECE diagnostic

The synthetic diagnostic YODA, developed in Paper I [88], closes the forward-modelling chain by coupling the kinetic distribution function to radiative transfer along the traced rays. For each channel and polarisation, the code imports the equilibrium and ray-path quantities — magnetic field strength B , parallel refractive index N_{\parallel} , and polarisation factors — evaluates the emission and absorption coefficients from the numerical f_e via Eqs. (4.8)–(4.11), and integrates the radiative transfer equation (4.4) to obtain the specific intensity I_{ω} at the antenna, with the boundary condition $I_{\omega}(s_0) = 0$ at the viewing dump.

A distinguishing feature of YODA relative to earlier synthetic ECE codes such as SPECE [89] is that the emission and absorption integrals are evaluated on the full numerical distribution function, rather than on parametric approximations such as sums of drifting Maxwellians. This

generality is important because the resonant interaction and quasilinear diffusion sculpt f_e in ways that analytic approximations cannot realistically reproduce [88].

Beyond enabling quantitative comparison with experiment, the forward model preserves a transparent link to the underlying phase-space dynamics. VECE channels with different frequency down-shifts probe different parts of the suprathermal tail, making it possible to constrain the energy dependence of \tilde{f} and to disentangle the respective roles of ohmic acceleration, quasilinear diffusion, and radial transport.

Chapter 5

Summary and outlook

This thesis addresses two complementary aspects of non-thermal electron dynamics in tokamak plasmas: the characterisation of suprathermal electrons as primary runaway seeds, and the generation and mitigation of runaway electrons during disruptions. While these may seem distinct, they are connected by the fact that a reliable estimation of the seed population is essential for predicting the runaway current that develops during a disruption. The four papers span a range from experimental diagnostics on present-day devices, where the YODA framework is developed and applied to study runaway seeds in existing tokamaks such as TCV, to predictive modelling of ITER disruption scenarios, where the sensitivity of the final runaway current to the seed is quantified. Below, a brief summary of each paper is given, followed by an outlook on future research directions.

5.1 Paper I – Experimental and numerical investigation of suprathermal electron dynamics using vertical electron cyclotron emission

This paper presents the development and application of YODA, a new synthetic electron cyclotron emission (ECE) diagnostic framework capable of computing ECE spectra from arbitrary numerical electron distribution functions. The code solves the radiation transport equation along ray paths obtained from the C3PO ray-tracing code, accounting for emission and reabsorption at multiple harmonics. YODA is validated against the established synthetic ECE code SPECE in a thermal TCV discharge, showing excellent agreement.

The framework is then applied to two TCV discharges featuring suprather-

mal electrons produced by electron cyclotron current drive. The electron distribution functions are obtained from the 3D bounce-averaged relativistic Fokker–Planck solver LUKE. In the first discharge (#72 644), where the EC launcher angle is varied in discrete steps, time-asymptotic LUKE simulations coupled with YODA reproduce the characteristic staircase pattern observed in the vertical ECE (VECE) signal. The analysis reveals that the 109 GHz channel (sensitive to ~ 20 keV electrons near the thermal bulk) responds to each angular step, whereas the 96 GHz channel (sensitive to ~ 90 keV electrons) only rises appreciably at larger toroidal angles, reflecting the slower build-up of the high-energy tail. In the second discharge (#73 217), time-dependent Fokker–Planck simulations capture the rapid decay of the VECE signal when EC heating is switched off, as well as a thermal ECE peak during the subsequent magnetic field ramp-down that enables an absolute calibration of the diagnostic. The simulations also demonstrate that both the ohmic electric field and the EC wave interaction must be included to accurately model the suprathreshold population.

5.2 Paper II – Runaway electron generation in ITER mitigated disruptions with improved physics models

This paper presents a comprehensive assessment of RE generation in SPI-mitigated ITER disruptions using an extended version of DREAM that incorporates four new physics models: (i) the reduced scrape-off loss model from Paper III; (ii) the plasmoid drift model developed in Ref. [63]; (iii) an adaptive hyper-resistive transport model; and (iv) an updated Compton scattering seed for the ITER tungsten first-wall design. The study covers 15 MA L-mode (H26) and H-mode (DTHmode24) scenarios, as well as a 7.5 MA H-mode case (DD7.5MA).

A systematic parameter scan confirms that complete avoidance of a multi-megaampere RE beam requires simultaneously: a pre-thermal-quench duration long enough to thermalise the hot-tail seed, sufficient hydrogen assimilation with limited neon content, and a representative seed current comparable to that of a single relativistic electron. In L-mode H26, staggered and low-neon single injections satisfy these conditions; in the DT H-mode scenario, however, the Compton seed prevents full RE suppression unless nuclear sources are artificially disabled.

Among the new models, VDE scrape-off losses play a decisive role in near-mitigated scenarios, where they can make the difference between a residual current and a multi-MA beam. Hyper-resistive current-profile re-

laxation reduces the available poloidal flux variation and thereby decreases the potential avalanche gain by up to four orders of magnitude. Rechester–Rosenbluth RE transport during the current quench reveals a clear threshold: for $\delta B/B \gtrsim 4 \times 10^{-4}$, the RE current is suppressed below 150 kA in all cases studied. A proposed two-stage SPI scheme for DT H-mode operation yields a representative RE current of only 0.24 MA, demonstrating a viable route to tolerable RE levels even with nuclear seeds.

5.3 Paper III – Reduced modelling of scrape-off losses of runaway electrons during tokamak disruptions

This paper introduces a reduced model for RE losses caused by the vertical displacement of the plasma during the current quench. The model exploits the observation, from 2D axisymmetric JOREK simulations, that the poloidal magnetic flux at the instantaneous last closed flux surface (LCFS) remains approximately constant during the current quench. This allows one to track the inward motion of the LCFS in DREAM’s otherwise fixed flux-surface geometry: flux surfaces whose self-consistent poloidal flux exceeds its initial edge value are deemed closed, while the region outside is treated as open and unable to sustain an RE current. A loss term removes REs from the open region on a time scale much shorter than any other relevant dynamics.

The model is validated against JOREK results, reproducing the avalanche gain to high accuracy. It is then applied to two SPI-mitigated ITER scenarios at 15 MA. In a non-nuclear L-mode case with relatively poor mitigation performance, the scrape-off losses reduce the RE plateau current only moderately, from 6.6 MA to 6.2 MA, because the RE current grows fast enough to halt the inward motion of the LCFS before the core flux surfaces are opened. In contrast, a DT H-mode scenario with a staggered injection and more favourable conditions produces a slower RE build-up, enabling all flux surfaces to be scraped off before a macroscopic RE beam forms; the maximum RE current is reduced to 250 kA. This sharp transition between the two regimes arises from a positive feedback loop: slower RE generation allows faster poloidal flux decay, which accelerates scrape-off and further suppresses RE growth. The results also reveal a significant sensitivity to the current-profile relaxation during the thermal quench, modelled here through hyper-resistivity.

5.4 Paper IV – Simulation of shattered pellet injections with plasmoid drifts in ASDEX Upgrade and ITER

This paper describes the implementation of a semi-analytical plasmoid drift model in DREAM and its application to SPI experiments at ASDEX Upgrade (AUG) and to ITER disruption scenarios. The model, based on the formulation in Vallhagen *et al.* (2023), calculates the radial displacement of the ablation cloud as a function of the cloud half-width Δy , temperature, composition, and background plasma conditions.

The model is first benchmarked against two AUG discharges. For a pure deuterium injection (#40 743), the experimental density profile measured by Thomson scattering is well reproduced with $\Delta y = 8.7$ mm, a value consistent with previous measurements and simulations. For a neon-doped injection (#40 732), the drift is strongly suppressed by radiative cooling of the ablation cloud, and the simulated density rise agrees with the experiment when material drifting past the LCFS is assumed to be recycled rather than lost. The results confirm that pure hydrogenic pellets are far more susceptible to drift losses than neon-doped pellets, owing to the higher plasmoid temperatures and pressures in the former case.

In ITER, the drift can reduce the assimilation of a pure deuterium first-stage pellet from 94% to as little as 4%, depending on the assumed Δy . Adding a small neon fraction ($\sim 0.3\%$) to the first pellet is found to largely restore the assimilation by cooling the ablation cloud and suppressing the drift. The resulting RE currents are, however, only moderately affected by the drift in the scenarios studied, because reduced assimilation of the first pellet is partly compensated by increased assimilation of the second, neon-doped pellet. A stronger sensitivity is observed when the thermal quench is fast and early, pointing to the possibility that it could become more important when RE scrape-off losses associated with the vertical plasma motion are taken into account.

5.5 Outlook

Several directions for future work emerge from this thesis. On the synthetic diagnostic side, the LUKE-YODA framework developed in Paper I opens the possibility of using VECE data, in combination with hard x-ray spectroscopy, to more tightly constrain the suprathreshold electron distribution function. Embedding YODA in a Bayesian inference framework would enable a systematic inversion of combined VECE and hard x-ray data to

recover the suprathreshold distribution function $f(p, \xi)$ together with quantified uncertainties.

On the disruption modelling side, the results of **Papers II–IV** identify several areas where further refinement is needed. The scrape-off model relies on the assumption of constant poloidal flux at the LCFS, which is completely violated during the RE plateau phase; relaxing this assumption would require coupling to a full MHD equilibrium evolution, which remains an open challenge. The plasmoid drift model is highly sensitive to the cloud half-width Δy , and developing a predictive, self-consistent model for this parameter, accounting for striation, plasmoid–plasmoid interaction, and scrape-off-layer dynamics, would greatly reduce the uncertainty in material assimilation predictions and in the effect of the pellet rocket force.

Crucially, hot-tail generation, which dominates the runaway seed in poorly mitigated ITER disruptions, is treated in the present work through the Smith–Verwichte analytical model, which is a simplified fluid model derived from a kinetic treatment. In reality, the interplay between stochastic magnetic transport, momentum and pitch-angle dynamics introduces qualitative features that a fluid model cannot capture. Overcoming this limitation will require extending the reduced kinetic models in DREAM to SPI disruption scenarios, capturing the complex momentum-space dynamics of the hot tail without the computational cost of the full kinetic equation.

A second limitation concerns electron–electron thermalisation. Following impurity injection, a cold Maxwellian population forms and gradually merges with the pre-disruption bulk. In the reduced kinetic models of DREAM, this cold population serves as the background f_0 against which the hot electrons scatter, and is assumed to remain Maxwellian throughout the process. In the fully kinetic approach, the same assumption enters through the linearised Fokker–Planck operator: at low injected densities, the perturbation δf to the cold population may become comparable to f_0 itself, rendering the linearisation inaccurate. In this regime, a nonlinear electron–electron collision operator is needed to describe the kinetic evolution of the injected population correctly, with potential consequences for both the thermalisation rate and the distribution function near the runaway boundary. Incorporating such an operator into DREAM is therefore an important direction for future work.

Declaration of Use of Generative AI

In writing this thesis, I made use of generative artificial intelligence tools as a support instrument during the drafting and editing process. The tools were mainly used to improve the clarity, grammar, and structure of selected passages, to help rephrase technical explanations, and to refine the overall readability of the text. In some cases, they were also used to assist with code-related tasks, such as debugging, commenting, and organising scripts used in the analysis.

The scientific work presented in this thesis, including the modelling choices, interpretation of results, figures, simulations, and conclusions, is my own. Any AI-assisted output was critically reviewed, modified where necessary, and checked for consistency with the underlying physics, numerical results, and cited literature. Generative AI was therefore used as an editorial and technical aid, not as a substitute for scientific reasoning or authorship.

References

- [1] Energy Institute, “Statistical review of world energy 2024,” 2024. Key finding: global energy demand reached 620 EJ in 2023.
- [2] U.S. Energy Information Administration, “International energy outlook 2023,” 2023. Projected growth in energy demand and drivers.
- [3] IPCC, *Climate Change 2022: Mitigation of Climate Change*. Intergovernmental Panel on Climate Change, 2022. Sixth Assessment Report, Working Group III, Chapter 6. Challenges of 100% renewable energy systems.
- [4] J. Wesson and D. J. Campbell, *Tokamaks*, vol. 149 of *International Series of Monographs on Physics*. Oxford, UK: Oxford University Press, 4th ed., 2011.
- [5] ITER Organization, “ITER: The way to new energy,” 2019. Brochure. Rationale for fusion in meeting long-term energy demand; fusion advantages.
- [6] K. S. Krane, *Introductory Nuclear Physics*. New York: John Wiley & Sons, 1987.
- [7] I. E. Tamm and A. D. Sakharov, “The theory of a magnetic thermonuclear reactor,” *Zh. Eksp. Teor. Fiz.*, vol. 21, pp. 635–646, 1951.
- [8] I. Bandyopadhyay, V. Igochine, O. Sauter, *et al.*, “MHD, disruptions and control physics: Chapter 4 of the special issue: on the path to tokamak burning plasma operation,” *Nuclear Fusion*, vol. 65, p. 103001, Sept. 2025.
- [9] P. Helander and D. J. Sigmar, *Collisional Transport in Magnetized Plasmas*, vol. 4 of *Cambridge Monographs on Plasma Physics*. Cambridge University Press, 2002.

- [10] E. Hirvijoki, I. Pusztai, J. Decker, O. Embréus, A. Stahl, and T. Fülöp, “Radiation reaction induced non-monotonic features in runaway electron distributions,” *Journal of Plasma Physics*, vol. 81, July 2015.
- [11] O. Embreus, A. Stahl, and T. Fülöp, “On the relativistic large-angle electron collision operator for runaway avalanches in plasmas,” *Journal of Plasma Physics*, vol. 84, no. 1, p. 905840102, 2018.
- [12] L. Hesslow, O. Embréus, G. J. Wilkie, G. Papp, and T. Fülöp, “Effect of partially ionized impurities and radiation on the effective critical electric field for runaway generation,” *Plasma Physics and Controlled Fusion*, vol. 60, p. 074010, June 2018.
- [13] J. Connor and R. Hastie, “Relativistic limitations on runaway electrons,” *Nuclear Fusion*, vol. 15, p. 415–424, June 1975.
- [14] C. T. R. Wilson, “The acceleration of β -particles in strong electric fields such as those of thunderclouds,” *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 22, no. 4, pp. 534–538, 1925.
- [15] H. Dreicer, “Electron and ion runaway in a fully ionized gas. i,” *Physical Review*, vol. 115, no. 2, pp. 238–249, 1959.
- [16] H. Dreicer, “Electron and ion runaway in a fully ionized gas. ii,” *Physical Review*, vol. 117, no. 2, pp. 329–342, 1960.
- [17] A. V. Gurevich, “On the theory of runaway electrons,” *Soviet Physics JETP*, vol. 12, no. 5, pp. 904–912, 1961.
- [18] H. Knoepfel and D. A. Spong, “Runaway electrons in tokamak plasmas,” *Nuclear Fusion*, vol. 19, no. 6, pp. 785–829, 1979.
- [19] M. Hoppe, O. Embreus, and T. Fülöp, “DREAM: a fluid-kinetic framework for tokamak disruption runaway electron simulations,” *Computer Physics Communications*, vol. 268, p. 108098, 2021.
- [20] M. D. Kruskal and I. B. Bernstein, “Runaway electrons in an ideal lorentz plasma,” *Physics of Fluids*, vol. 7, p. 407, 1964.
- [21] H. M. Smith and E. Verwichte, “Hot tail runaway electron generation in tokamak disruptions,” *Physics of Plasmas*, vol. 15, no. 7, p. 072502, 2008.

- [22] I. Ekmark, M. Hoppe, T. Fülöp, P. Jansson, L. Antonsson, O. Vallhagen, and I. Pusztai, “Fluid and kinetic studies of tokamak disruptions using bayesian optimization,” *Journal of Plasma Physics*, vol. 90, May 2024.
- [23] J. R. Martín-Solís, A. Loarte, and M. Lehnen, “Formation and termination of runaway beams in ITER disruptions,” *Nuclear Fusion*, vol. 57, no. 6, p. 066025, 2017.
- [24] O. Vallhagen, O. Embreus, I. Pusztai, L. Hesslow, and T. Fülöp, “Runaway dynamics in the DT phase of ITER operations in the presence of massive material injection,” *Journal of Plasma Physics*, vol. 86, no. 4, p. 475860401, 2020.
- [25] B. N. Breizman, P. Aleynikov, E. M. Hollmann, and M. Lehnen, “Physics of runaway electrons in tokamaks,” *Nuclear Fusion*, vol. 59, no. 8, p. 083001, 2019.
- [26] Y. A. Sokolov, ““Multiplication” of Accelerated Electrons in a Tokamak,” *JETP Letters*, vol. 29, pp. 218–221, 1979.
- [27] R. Jayakumar, H. H. Fleischmann, and S. J. Zweben, “Collisional avalanche exponentiation of runaway electrons in electrified plasmas,” *Physics Letters A*, vol. 172, no. 6, pp. 447–451, 1993.
- [28] M. N. Rosenbluth and S. V. Putvinski, “Theory for avalanche of runaway electrons in tokamaks,” *Nuclear Fusion*, vol. 37, no. 10, pp. 1355–1362, 1997.
- [29] L. Hesslow, O. Embreus, M. Hoppe, T. C. DuBois, G. Papp, M. Rahm, and T. Fülöp, “Generalized collision operator for fast electrons interacting with partially ionized impurities,” *Journal of Plasma Physics*, vol. 84, no. 6, p. 905840605, 2018.
- [30] L. Hesslow, O. Embréus, A. Stahl, T. C. DuBois, G. Papp, S. L. Newton, and T. Fülöp, “Effect of partially screened nuclei on fast-electron dynamics,” *Phys. Rev. Lett.*, vol. 118, p. 255001, Jun 2017.
- [31] L. Hesslow, O. Embreus, O. Vallhagen, and T. Fülöp, “Influence of massive material injection on avalanche runaway generation during tokamak disruptions,” *Nuclear Fusion*, vol. 59, no. 8, p. 084004, 2019.

- [32] A. B. Rechester and M. N. Rosenbluth, “Electron heat transport in a tokamak with destroyed magnetic surfaces,” *Physical Review Letters*, vol. 40, no. 1, pp. 38–41, 1978.
- [33] T. Hauff and F. Jenko, “Runaway electron transport via tokamak microturbulence,” *Physics of Plasmas*, vol. 16, no. 10, p. 102308, 2009.
- [34] P. Svensson, O. Embreus, S. L. Newton, K. Särkimäki, O. Vallhagen, and T. Fülöp, “Effects of magnetic perturbations and radiation on the runaway avalanche,” *Journal of Plasma Physics*, vol. 87, no. 2, p. 905870211, 2021.
- [35] O. Vallhagen, *Disruption mitigation in tokamaks with massive material injection*. PhD thesis, Chalmers University of Technology, Göteborg, Sweden, 2025.
- [36] D. I. Kiramov and B. N. Breizman, “Model of vertical plasma motion during the current quench,” *Physics of Plasmas*, vol. 24, no. 10, p. 100702, 2017.
- [37] O. Vallhagen, L. Hanebring, T. Fülöp, M. Hoppe, L. Votta, and I. Pusztai, “Reduced modelling of scrape-off losses of runaway electrons during tokamak disruptions,” *Journal of Plasma Physics*, vol. 91, no. 3, p. E78, 2025.
- [38] C. Wang, E. Nardon, F. Artola, V. Bandaru, and M. Hoelzl, “The effect of vertical displacements on the runaway electron avalanche in iter mitigated disruptions,” *Nuclear Fusion*, vol. 65, no. 1, p. 016012, 2025.
- [39] E. M. Hollmann, P. B. Aleynikov, T. Fülöp, D. A. Humphreys, V. A. Izzo, M. Lehnen, V. E. Lukash, G. Papp, G. Pautasso, F. Saint-Laurent, and J. A. Snipes, “Status of research toward the ITER disruption mitigation system,” *Physics of Plasmas*, vol. 22, no. 2, 2015.
- [40] C. Reux, V. Plyusnin, B. Alper, D. Alves, B. Bazylev, E. Belonohy, A. Boboc, S. Brezinsek, I. Coffey, J. Decker, P. Drewelow, S. Devaux, P. de Vries, A. Fil, S. Gerasimov, L. Giacomelli, S. Jachmich, E. Khilkevitch, V. Kiptily, R. Koslowski, U. Kruezi, M. Lehnen, I. Lupelli, P. Lomas, A. Manzanares, A. Martin De Aguilera, G. Matthews, J. Mlynář, E. Nardon, E. Nilsson, C. Perez von Thun, V. Riccardo,

- F. Saint-Laurent, A. Shevelev, G. Sips, and C. Sozzi, “Runaway electron beam generation and mitigation during disruptions at JET-ILW,” *Nuclear Fusion*, vol. 55, p. 093013, Aug. 2015.
- [41] L. Baylor, S. Combs, C. Foust, T. Jernigan, S. Meitner, P. Parks, J. Caughman, D. Fehling, S. Maruyama, A. Qualls, D. Rasmussen, and C. Thomas, “Pellet fuelling, ELM pacing and disruption mitigation technology development for ITER,” *Nuclear Fusion*, vol. 49, p. 085013, July 2009.
- [42] N. Commaux, L. Baylor, T. Jernigan, E. Hollmann, P. Parks, D. Humphreys, J. Wesley, and J. Yu, “Demonstration of rapid shut-down using large shattered deuterium pellet injection in DIII-D,” *Nuclear Fusion*, vol. 50, p. 112001, Sept. 2010.
- [43] ITER Organization, “Disruption mitigation.” Web page, 2023. Accessed 2026-02-10.
- [44] M. Lehnen, “The ITER disruption mitigation strategy.” IAEA Technical Meeting on Plasma Disruptions and their Mitigation, 2020.
- [45] ITER Organization, “On the possible injection schemes with the ITER SPI system.” IAEA Technical Meeting on Plasma Disruptions and their Mitigation, 2020.
- [46] O. Vallhagen, I. Pusztai, M. Hoppe, S. L. Newton, and T. Fülöp, “Effect of two-stage shattered pellet injection on tokamak disruptions,” *Nuclear Fusion*, vol. 62, p. 112004, 2022.
- [47] S. Jachmich, U. Kruezi, M. Lehnen, M. Baruzzo, L. Baylor, D. Carnevale, D. Craven, N. Eidietis, O. Ficker, T. Gebhart, S. Gerasimov, J. Herfindal, E. Hollmann, A. Huber, P. Lomas, J. Lovell, A. Manzanares, M. Maslov, J. Mlynar, G. Pautasso, C. Paz-Soldan, A. Peacock, L. Piron, V. Plyusnin, M. Reinke, C. Reux, F. Rimini, U. Sheikh, D. Shiraki, S. Silburn, R. Sweeney, J. Wilson, P. Carvalho, and JET Contributors, “Shattered pellet injection experiments at JET in support of the ITER disruption mitigation system design,” *Nuclear Fusion*, vol. 62, p. 026012, Dec. 2021.
- [48] M. Dibon, P. de Marne, G. Papp, I. Vinyar, A. Lukin, S. Jachmich, U. Kruezi, A. Muir, V. Rohde, M. Lehnen, P. Heinrich, T. Peherstorfer, and D. Podymskii, “Design of the shattered pellet injection system

- for ASDEX Upgrade,” *Review of Scientific Instruments*, vol. 94, Apr. 2023.
- [49] S. Park, K. Lee, L. R. Baylor, S. J. Meitner, H. Lee, J. Song, T. E. Gebhart, S. Yun, J. Kim, K. Kim, K. Park, and S. Yoon, “Deployment of multiple shattered pellet injection systems in KSTAR,” *Fusion Engineering and Design*, vol. 154, p. 111535, 2020.
- [50] A. Fil, L. Henden, S. Newton, M. Hoppe, and O. Vallhagen, “Disruption runaway electron generation and mitigation in the Spherical Tokamak for Energy Production (STEP),” *Nuclear Fusion*, vol. 64, no. 10, p. 106049, 2024.
- [51] P. Zhu, L. Li, Y. Fang, Y. He, S. Wang, R. Han, Y. Liu, X. Wang, Y. Zhang, X. Zhang, Q. Yu, L. Hu, H. Wang, Y. Sun, L. Wei, W. Tang, T. Liu, Z. Wang, X. Yan, W. Huang, Y. Hou, X. Ji, S. Zeng, Z. Abdullah, Z. Chen, L. Zeng, H. Li, Z. Chen, Z. Wang, B. Rao, M. Zhang, Y. Ding, Y. Pan, and CFETR Physics Team, “MHD analysis on the physics design of CFETR baseline scenarios,” *Journal of Fusion Energy*, vol. 41, 2022.
- [52] N. Mott and E. Linfoot, *A Theory of Fragmentation*, pp. 207–225. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006.
- [53] P. Parks, “Modeling dynamic fracture of cryogenic pellets,” tech. rep., General Atomics, 2016. Technical Report GA–A28352.
- [54] T. E. Gebhart, L. R. Baylor, and S. J. Meitner, “Experimental pellet shatter thresholds and analysis of shatter tube ejecta for disruption mitigation cryogenic pellets,” *IEEE Transactions on Plasma Science*, vol. 48, no. 6, pp. 1598–1605, 2020.
- [55] N. J. Guth, O. Vallhagen, P. Helander, I. Pusztai, S. L. Newton, and T. Fülöp, “The pellet rocket effect in magnetic confinement fusion plasmas,” *Physical Review Letters*, vol. 134, p. 035101, 2025.
- [56] P. B. Parks and R. J. Turnbull, “Effect of transonic flow in the ablation cloud on the lifetime of a solid hydrogen pellet in a plasma,” *Physics of Fluids*, vol. 21, pp. 1735–1741, 1978.
- [57] B. Pégourié, “Review: Pellet injection experiments and modelling,” *Plasma Physics and Controlled Fusion*, vol. 49, pp. R87–R122, 2007.

- [58] A. K. Fontanilla and B. N. Breizman, “Heating and ablation of high- z cryogenic pellets in high temperature plasmas,” *Nuclear Fusion*, vol. 59, p. 096033, July 2019.
- [59] P. B. Parks, “A theoretical model for the penetration of a Shattered-Pellet debris plume.” Presented at the Theory and Simulation of Disruptions Workshop, Princeton Plasma Physics Laboratory, Princeton, NJ, USA, 17–19 July 2017, July 2017. Presentation slides.
- [60] H. W. Müller, K. Büchl, M. Kaufmann, P. T. Lang, R. S. Lang, A. Lorenz, M. Maraschek, V. Mertens, J. Neuhauser, and ASDEX Upgrade Team, “High- β plasmoid drift during pellet injection into tokamaks,” *Physical Review Letters*, vol. 83, pp. 2199–2202, 1999.
- [61] P. B. Parks, W. D. Sessions, and L. R. Baylor, “Radial displacement of pellet ablation material in tokamaks due to the grad- b effect,” *Physics of Plasmas*, vol. 7, pp. 1968–1975, 2000.
- [62] V. Rozhansky, I. Senichenkov, I. Y. Veselova, and R. Schneider, “Mass deposition after pellet injection into a tokamak,” *Plasma Physics and Controlled Fusion*, vol. 46, pp. 575–591, 2004.
- [63] O. Vallhagen, I. Pusztai, P. Helander, S. Newton, and T. Fülöp, “Drift of ablated material after pellet injection in a tokamak,” *Journal of Plasma Physics*, vol. 89, no. 3, p. 905890306, 2023.
- [64] O. Vallhagen, L. Antonsson, P. Haldestam, G. Papp, P. Heinrich, A. Patel, M. Hoppe, L. Votta, ASDEX Upgrade Team, and EUROfusion Tokamak Exploitation Team, “Simulation of shattered pellet injections with plasmoid drifts in ASDEX Upgrade and ITER,” *Plasma Physics and Controlled Fusion*, vol. 67, p. 105034, 2025.
- [65] L. Votta, F. J. Artola, E. Nardon, O. Vallhagen, and M. Hoppe, “Runaway electron generation in ITER mitigated disruptions with improved physics models,” Submitted to Nuclear Fusion preprint arXiv:2602.22177, 2026.
- [66] E. Nardon, D. Hu, M. Hoelzl, D. Bonfiglio, JET Contributors, and JOEKE Team, “Fast plasma dilution in ITER with pure deuterium shattered pellet injection,” *Nuclear Fusion*, vol. 60, p. 126040, 2020.

- [67] D. Hu, E. Nardon, M. Lehnen, G. T. A. Huijsmans, and D. C. van Vugt, “3D non-linear MHD simulation of the MHD response and density increase as a result of shattered pellet injection,” *Nuclear Fusion*, vol. 58, p. 126025, 2018.
- [68] H. P. Summers, *The ADAS User Manual*. ADAS, version 2.6 ed., 2004.
- [69] D. I. Kiramov, M. Lehnen, R. Khayrutdinov, and V. Lukash, “ITER disruption simulations with improved power balance in the halo region,” in *Proceedings of the 43rd EPS Conference on Plasma Physics*, 2016. Preprint.
- [70] A. H. Boozer, “Ohm’s law for mean magnetic fields,” *Journal of Plasma Physics*, vol. 35, pp. 133–139, 1986.
- [71] A. H. Boozer, “Pivotal issues on relativistic electrons in ITER,” *Nuclear Fusion*, vol. 58, p. 036006, 2018.
- [72] A. Lvovskiy, C. Paz-Soldan, N. Eidietis, A. Dal Molin, M. Nocente, C. Cooper, D. Rigamonti, M. Tardocchi, and D. Taussig, “Upgrades to the gamma ray imager on DIII-D enabling access to high flux hard x-ray measurements during the runaway electron plateau phase,” *Review of Scientific Instruments*, vol. 93, Nov. 2022.
- [73] A. Shevelev, E. Khilkevitch, S. Lashkul, V. Rozhdestvensky, A. Al-tukhov, I. Chugunov, D. Doinikov, L. Esipov, D. Gin, M. Iliasova, V. Naidenov, N. Nersesyan, I. Polunovsky, A. Sidorov, and V. Kip-tily, “High performance gamma-ray spectrometer for runaway electron studies on the ft-2 tokamak,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 830, p. 102–108, Sept. 2016.
- [74] A. Dal Molin, L. Fumagalli, M. Nocente, D. Rigamonti, M. Tardocchi, L. Giacomelli, E. Panontin, A. Lvovskiy, C. Paz-Soldan, N. W. Eidietis, and G. Gorini, “Novel compact hard x-ray spectrometer with mcps counting rate capabilities for runaway electron measurements on DIII-D,” *Review of Scientific Instruments*, vol. 92, Apr. 2021.
- [75] R. Jaspers, N. J. Lopes Cardozo, A. J. H. Donné, H. L. M. Withershoven, and K. H. Finken, “A synchrotron radiation diagnostic to observe relativistic runaway electrons in a tokamak plasma,” *Rev. Sci. Instrum.*, vol. 72, pp. 466–470, 2001.

- [76] M. Hoppe, O. Embreus, R. A. Tinguely, R. S. Granetz, A. Stahl, and T. Fülöp, “SOFT: a synthetic synchrotron diagnostic for runaway electrons,” *Nucl. Fusion*, vol. 58, p. 026032, 2018.
- [77] C. Paz-Soldan, C. M. Cooper, P. Aleynikov, D. C. Pace, N. W. Eidietis, D. P. Brennan, R. S. Granetz, E. M. Hollmann, C. Liu, A. Lvovskiy, R. A. Moyer, and D. Shiraki, “Spatiotemporal evolution of runaway electron momentum distributions in tokamaks,” *Phys. Rev. Lett.*, vol. 118, p. 255002, 2017.
- [78] T. C. Luce, P. C. Efthimion, N. J. Fisch, R. E. Bell, and J. E. Stevens, “Vertical viewing of ECE from non-thermal distributions,” *AIP Conf. Proc.*, vol. 159, pp. 167–170, 1987.
- [79] K. Kato, *Diagnosis of mildly relativistic electron velocity distributions by electron cyclotron emission in the Alcator C Tokamak*. PhD thesis, Massachusetts Institute of Technology, 1986.
- [80] G. Giruzzi, J.-L. Ségui, A.-L. Pecquet, and C. Gil, “Observation of the $m = 1$ mode by microwave transmission measurements in Tore Supra,” *Nuclear Fusion*, vol. 31, p. 2158–2162, Nov. 1991.
- [81] G. Giruzzi, J. L. Ségui, T. D. de Wit, Y. Michelot, Y. Peysson, D. Moreau, and M. Talvard, “Measurement of the time constants of fast electron distributions in the Tore Supra tokamak,” *Physical Review Letters*, vol. 74, p. 550–553, Jan. 1995.
- [82] S. J. Janz, *Analysis of nonthermal electron cyclotron emission during electron cyclotron current drive experiments on the DIII-D Tokamak*. PhD thesis, University of Maryland, 1992.
- [83] A. Tema Biwolé, L. Porte, S. Coda, and A. Fasoli, “Vertical electron cyclotron emission diagnostic on the tokamak à configuration variable,” *Review of Scientific Instruments*, vol. 94, Oct. 2023.
- [84] A. Tema Biwole, L. Porte, A. Fasoli, L. Figini, J. Decker, M. Hoppe, J. Cazabonne, L. Votta, A. Simonetto, and S. Coda, “Cross-calibration and first vertical ECE measurement of electron energy distribution in the TCV tokamak,” *Plasma Physics and Controlled Fusion*, vol. 66, p. 125010, Nov. 2024.

- [85] P. Blanchard, *Études du rayonnement suprathermique émis lors du chauffage cyclotronique électronique du plasma du tokamak TCV*. PhD thesis, École Polytechnique Fédérale de Lausanne, 2002.
- [86] I. Klimanov, *Reconstruction of the electron distribution function during ECRH/ECCD and magnetic reconnection events in a tokamak plasma*. PhD thesis, École Polytechnique Fédérale de Lausanne, 2005.
- [87] A. Tema Biwolé, L. Porte, A. Fasoli, A. Simonetto, and O. D’Arcangelo, “Performance of a high vacuum, high temperature compatible millimeter-range viewing dump for the vertical ECE experiment on TCV,” *Fusion Engineering and Design*, vol. 162, p. 112079, 2021.
- [88] L. Votta, M. Hoppe, J. Decker, E. Devlaminck, A. S. Tema Biwolé, L. Porte, J. Cazabonne, Y. Savoye-Peysson, and TCV Team, “Experimental and numerical investigation of suprathermal electron dynamics using vertical electron cyclotron emission,” *Plasma Phys. Control. Fusion*, vol. 68, p. 015029, 2026.
- [89] L. Figini, *Electron Cyclotron Emission in Tokamaks: development of a new modeling tool for data validation, analysis and predictions*. PhD thesis, Università degli studi di Milano, 2009.
- [90] S. Denk, *Study of non-thermal energy distribution functions by modeling of electron cyclotron emission from the fusion plasma of the ASDEX Upgrade Tokamak*. PhD thesis, Technical University of Munich, 2019.
- [91] G. Bekefi, *Radiation Processes in Plasmas*. Wiley, 1966.
- [92] M. Bornatici, R. Cano, O. D. Barbieri, and F. Engelmann, “Electron cyclotron emission and absorption in fusion plasmas,” *Nucl. Fusion*, vol. 23, p. 1153, 1983.
- [93] Y. Peysson, J. Decker, and L. Morini, “A versatile ray-tracing code for studying RF wave propagation in toroidal magnetized plasmas,” *Plasma Phys. Control. Fusion*, vol. 54, p. 045003, 2012.
- [94] P. Bonoli, “Linear theory of lower hybrid heating,” *IEEE Trans. Plasma Sci.*, vol. 12, pp. 95–107, 1984.

- [95] J. P. Bizarro, Y. Peysson, P. T. Bonoli, J. Carrasco, T. Dudok de Wit, V. Fuchs, G. T. Hoang, X. Litaudon, D. Moreau, C. Pocheau, and I. P. Shkarofsky, “On self-consistent ray-tracing and Fokker–Planck modeling of the lower-hybrid current drive in tokamaks,” *Physics of Fluids B: Plasma Physics*, vol. 5, no. 9, pp. 3276–3283, 1993.
- [96] Y. Peysson, R. Arslanbekov, V. Basiuk, J. Carrasco, X. Litaudon, D. Moreau, and J. P. Bizarro, “Full wave and ray tracing analysis of lower hybrid wave propagation,” *Phys. Plasmas*, vol. 3, pp. 3668–3688, 1996.
- [97] F. Imbeaux and Y. Peysson, “Ray-tracing and Fokker–Planck modelling of the effect of plasma current on the propagation and absorption of lower hybrid waves,” *Plasma Physics and Controlled Fusion*, vol. 47, p. 2041–2065, Oct. 2005.
- [98] F. Albajar, N. Bertelli, M. Bornatici, and F. Engelmann, “Electron-cyclotron absorption in high-temperature plasmas: quasi-exact analytical evaluation and comparative numerical analysis,” *Plasma Physics and Controlled Fusion*, vol. 49, p. 15–29, Nov. 2006.
- [99] U. Belotti, M. Bornatici, and F. Engelmann, “Electron cyclotron radiation theory,” *Riv. Nuovo Cimento*, vol. 20, pp. 1–67, 1997.
- [100] J. Decker and Y. Peysson, “DKE: a fast numerical solver for the 3D drift kinetic equation,” Tech. Rep. EUR-CEA-FC-1736, Euratom-CEA, 2004.

