





Disruption Mitigation in Tokamaks with Shattered Pellet Injection

Thesis for the degree of Master of Science in Physics

OSKAR VALLHAGEN

Department of Physics CHALMERS UNIVERSITY OF TECHNOLOGY Gothenburg, Sweden 2021

Master's thesis 2021:NN

Disruption Mitigation in Tokamaks with Shattered Pellet Injection

Oskar Vallhagen



Department of Physics Division of Subatomic, High Energy and Plasma Physics CHALMERS UNIVERSITY OF TECHNOLOGY Gothenburg, Sweden 2021 Disruption Mitigation in Tokamaks with Shattered Pellet Injection Oskar Vallhagen

© Oskar Vallhagen, 2021.

Supervisors:	Istvan Pusztai, Department of Physics
	Sarah Newton, Department of Physics;
	and Culham Centre for Fusion Energy, UK
Examiner:	Tünde Fülöp, Department of Physics

Department of Physics Division of Subatomic, High Energy and Plasma Physics Chalmers University of Technology SE-412 96 Gothenburg Telephone +46 31 772 1000

Cover: Contour plots of the deuterium density at different times during a shattered pellet injection in a circular cross-section tokamak plasma, simulated with DREAM. The color scale ranges from 10^{20} m^{-3} (black) to $7 \cdot 10^{21} \text{ m}^{-3}$ (white). The pellet shard positions are marked by cyan dots, whose sizes are proportional to the shard radii.

Typeset in $\[\]ATEX$ Printed by Chalmers Reproservice Gothenburg, Sweden 2021 Disruption Mitigation in Tokamaks with Shattered Pellet Injection

Oskar Vallhagen Department of Physics Chalmers University of Technology

Abstract

The sudden loss of confinement of the energy content of fusion plasmas in off-normal events, called disruptions, are among the most severe threats to the future of fusion energy based on the tokamak design. An efficient disruption mitigation system will therefore be of utmost importance for future large, high-current devices such as ITER. The potentially greatest threat to be mitigated is posed by currents carried by highly energetic electrons, called runaway electrons, which may cause severe damage upon wall impact. The disruption mitigation system must also ensure a sufficiently homogeneous deposition of the thermal energy on the plasma-facing components, and avoid excessive forces on the machine due to currents flowing in the surrounding structures. The currently envisaged mitigation method is to make a massive material injection when an emerging disruption is detected, attempting to better control the plasma cooling and energy dissipation. In this work, we perform numerical simulations assessing the performance of the most up to date mitigation schemes based on shattered pellet injections in an ITER-like setting, with a particular focus on the generation of runaway electrons. The main mitigation scheme investigated is a two-stage shattered pellet injection, with a diluting deuterium injection followed by a neon injection aiming to radiatively dissipate the plasma energy content.

Our studies indicate that the diluting deuterium injection can efficiently reduce the runaway generation due to the hot-tail mechanism, by allowing for an intermediate equilibration of the superthermal electron population between the injections. The fraction of the initial thermal energy content conducted to the plasma-facing components is also reduced compared to a single-stage injection with the same composition, reducing the localised heat loads. During non-nuclear operation, the maximum runaway current was found to be reduced to acceptable levels with realistic two-stage injection parameters. On the other hand, during nuclear operation, the unavoidable runaway seed from tritium decay and compton scattering was found to be amplified to several mega-amperes by the avalanche mechanism for all investigated injection parameters. The reason is that the intense cooling from the injected material leads to a high induced electric field and a substantial recombination, resulting in an enhanced avalanche multiplication.

Keywords: fusion plasma, disruption mitigation, shattered pellet injection, runaway electron

Acknowledgements

I would like to express my sincerest gratitude towards all the people who have supported me in my early scientific career, leading up to the finishing of this project. First of all, I want to thank my examiner and mentor Tünde Fülöp, not only for being my examiner on this project, but primarily for giving me the chance to work with this wonderful group for the last three years. Her generous guidance and support has really given me an ideal start of my scientific career. During this project, as well as others, I have also received invaluable support from my supervisors István Pusztai and Sarah Newton, which greatly increased the quality of my work and learning far beyond what would have been possible without them. I also want to thank my unofficial supervisors Mathias Hoppe and Ola Embreus, for their enthusiastic development of the DREAM code, and for listening to all my thoughts and questions with a similar enthusiasm and expertise. My learning curve within the field of plasma physics has also been greatly helped by several additional former and present colleagues, of which I especially want to thank Linnea Hesslow and Gergely Papp. Thank you for the great cooperation over the years, and for helping me get such a good introduction to the field.

Moving outside the group, I want to thank all my fellow students and friends, who have made my time at Chalmers such an educational and enjoyable experience. A special thanks goes to Julia Järlebark, for the countless hours of hard and rewarding work, exam panic, and good fun for the last six years, and to Erik Svensson for all the philosophical conversations, spontaneous singing, and moral support in stressful situations. Finally, I want to thank my family, for their company, their support, for coping with me working from home during the pandemic, and for forcing me to take breaks once in a while by shooting me in the head with various sorts of toy guns.

Oskar Vallhagen, Alafors, March 2021

Contents

1	Intr	oduction		
	1.1	Magnetically confined fusion		
	1.2	Tokamak disruptions		
		1.2.1 Runaway electrons		
		1.2.2 Mitigation techniques		
	1.3	Thesis outline		
2	Rev	view of the underlying plasma physics 7		
	2.1	Definition and basic properties		
	2.2	Magnetically confined plasmas		
		2.2.1 Charged particle motions in electromagnetic fields 8		
		2.2.2 The tokamak design		
	2.3	Collisions in a plasma		
		2.3.1 Coulomb collisions		
		2.3.2 Conductivity		
		2.3.3 The runaway phenomenon		
	2.4	Theoretical models of plasmas		
		2.4.1 Kinetic models		
		2.4.1.1 The Coulomb collision operator		
		2.4.2 Fluid models		
		2.4.3 Electromagnetic fields		
3	Tokamak Disruptions			
	3.1	General features		
	3.2	Vessel loads and mitigation requirements		
	3.3	Runaway electrons		
	3.4	Disruption mitigation strategies		
		3.4.1 Massive material injection		
		$3.4.1.1$ Gas injection $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots 31$		
		$3.4.1.2$ Pellet injection $\ldots \ldots \ldots \ldots \ldots \ldots \ldots 32$		
		3.4.1.3 Shattered pellet injection		
		3.4.2 Injection schemes		
		3.4.2.1 Mixed deuterium-impurity injection		
		3.4.2.2 Multiple-stage injection		
	3.5	Pellet injection physics		
		3.5.1 Pellet ablation		

	3.5.2	Pellet material deposition	•	•				•	36
Disr	Disruption model								
4.1 Shattered pellet injection									39
	4.1.1	Shattering							41
	4.1.2	Ablation							42
	4.1.3	Material deposition							43
4.2	Particl	e and energy balance							44
4.3	Electri	c field evolution							47
4.4 Runaway generation									48
	4.4.1	Kinetic equation							48
	4.4.2	Fluid runaway sources							50
4.5	Numer	ical implementation and settings \ldots \ldots \ldots		•					52
Sim	ulation	s of two-stage SPI injections							57
5.1	First s	tage deuterium injections	_	_					58
5.2	Second	stage neon injections			·	·			62
0.2	5 2 1	Single vs two-stage injection dynamics	•	•	•	·	·	•	63
	522	Radiative vs transport losses	•	•	•	•	•	•	68
	523	Hot-tail generation	·	•	•	•	•	•	69
5.3	Currer	t auench simulations	·	•	•	•	•	•	72
0.0	5.3.1	Current quench and runaway dynamics	·	•	·	·	•	·	76
	5.3.2	Qualitative interpretation based on energy balance	•	•	•	•	•	•	79
Con	aludin	n nomenles							ວງ
6 1	Summ	g remarks							0 0 02
0.1 6 0	Outloo	ary of results	·	•	·	·	·	·	00 05
0.2	Outloc	9K	•	•	•	•	·	•	60
A Discretisation of SPI deposition kernels									Ι
Opacity for deuterium radiation					III				
C Impact of the deposition process features						VII			
	Disr 4.1 4.2 4.3 4.4 4.5 Sim 5.1 5.2 5.3 5.3 Con 6.1 6.2 Disc Opa Imp	3.5.2 Disruption 4.1 Shatte 4.1.1 4.1.2 4.1.3 4.2 Particl 4.3 Electri 4.4 Runaw 4.4.1 4.4.2 4.5 Numer 5.1 First s 5.2 Second 5.2.1 5.2.2 5.2.3 5.3 Curren 5.3.1 5.3.2 Concludin 6.1 Summ 6.2 Outloo Discretisat Opacity fo Impact of	3.5.2 Pellet material deposition Disruption model 4.1 Shattered pellet injection 4.1.1 Shattering 4.1.2 Ablation 4.1.3 Material deposition 4.1.4 Ablation 4.1.2 Ablation 4.1.3 Material deposition 4.1.4 Ablation 4.1.5 Particle and energy balance 4.1 Runaway generation 4.4 Runaway generation 4.4.1 Kinetic equation 4.4.2 Fluid runaway sources 4.5 Numerical implementation and settings 5.1 First stage deuterium injections 5.1 First stage deuterium injections 5.2.1 Single vs two-stage injection dynamics 5.2.2 Radiative vs transport losses 5.2.3 Hot-tail generation 5.3.1 Current quench simulations 5.3.2 Qualitative interpretation based on energy balance Concluding remarks 6.1 6.1 Summary of results 6.2 Outlook Outlook Impact of the deposit	3.5.2 Pellet material deposition Disruption model 4.1 Shattered pellet injection 4.1.1 Shattering 4.1.2 Ablation 4.1.3 Material deposition 4.1.4 Ablation 4.1.2 Ablation 4.1.3 Material deposition 4.1.4 Ablation 4.1.5 Particle and energy balance 4.1.6 Kinetic field evolution 4.1 Kinetic equation 4.4.1 Kinetic equation 4.4.2 Fluid runaway sources 4.5 Numerical implementation and settings 5.1 First stage deuterium injections 5.1 First stage deuterium injections 5.2 Second stage neon injections 5.2.1 Single vs two-stage injection dynamics 5.2.2 Radiative vs transport losses 5.3.1 Current quench and runaway dynamics 5.3.2 Qualitative interpretation based on energy balance 5.3.2 Qualitative interpretation based on energy balance 6.2 Outlook 6.3 Summary of results	3.5.2 Pellet material deposition Disruption model 4.1 Shattered pellet injection 4.1.1 Shattering 4.1.2 Ablation 4.1.3 Material deposition 4.1 Shattering 4.1.2 Ablation 4.1.3 Material deposition 4.1 Shattering 4.1.3 Material deposition 4.1 Ablation 4.1 Runaway generation 4.4 Runaway generation 4.4.1 Kinetic equation 4.4.2 Fluid runaway sources 4.5 Numerical implementation and settings 5.1 First stage deuterium injections 5.1 First stage deuterium injections 5.2 Second stage neon injections 5.2.1 Single vs two-stage injection dynamics 5.2.2 Radiative vs transport losses 5.3.1 Current quench simulations 5.3.1 Current quench and runaway dynamics 5.3.2 Qualitative interpretation based on energy balance 5.3.2 Qualitative interpretation based on energy balance	3.5.2 Pellet material deposition Disruption model 4.1 Shattered pellet injection 4.1.1 Shattering 4.1.2 Ablation 4.1.3 Material deposition 4.1 Shattering 4.1.2 Ablation 4.1.3 Material deposition 4.1 Shattering 4.1.3 Material deposition 4.1 Ablation 4.2 Particle and energy balance 4.3 Electric field evolution 4.4 Runaway generation 4.4.1 Kinetic equation 4.4.2 Fluid runaway sources 4.5 Numerical implementation and settings 5.1 First stage deuterium injections 5.1 First stage deuterium injections 5.2 Second stage neon injections 5.2.1 Single vs two-stage injection dynamics 5.2.2 Radiative vs transport losses 5.3.1 Current quench simulations 5.3.1 Current quench and runaway dynamics 5.3.2 Qualitative interpretation based on energy balance 5.3.2	3.5.2 Pellet material deposition Disruption model 4.1 Shattered pellet injection 4.1.1 Shattering 4.1.2 Ablation 4.1.3 Material deposition 4.1 Anterial deposition 4.1.3 Material deposition 4.1 Ablation 4.1.3 Material deposition 4.1 Ablation 4.1.4 Ablation 4.1 Particle and energy balance 4.3 Electric field evolution 4.4 Runaway generation 4.4.1 Kinetic equation 4.4.2 Fluid runaway sources 4.5 Numerical implementation and settings 5.1 First stage deuterium injections 5.2 Second stage neon injections 5.2.1 Single vs two-stage injection dynamics 5.2.2 Radiative vs transport losses 5.3.1 Current quench and runaway dynamics 5.3.2 Qualitative interpretation based on energy balance 5.3.2 Qualitative interpretation based on energy balance 5.3.2 Outlook Co	3.5.2 Pellet material deposition Disruption model 4.1 Shattered pellet injection 4.1.1 Shattering 4.1.2 Ablation 4.1.3 Material deposition 4.1.4 Ablation 4.1.5 Material deposition 4.1.6 Anterial deposition 4.1.7 Ablation 4.1.8 Material deposition 4.1 Material deposition 4.2 Particle and energy balance 4.3 Electric field evolution 4.4 Runaway generation 4.4.1 Kinetic equation 4.4.2 Fluid runaway sources 4.5 Numerical implementation and settings 5.1 First stage deuterium injections 5.2 Second stage neon injections 5.2.1 Single vs two-stage injection dynamics 5.3.1 Current quench simulations 5.3.2 Qualitative interpretation based on energy balance 5.3.2 Qualitative interpretation based on energy balance 5.3.2 Qualitative interpretation based on energy balance 6.2 Outlook </td <td>3.5.2 Pellet material deposition Disruption model 4.1 Shattered pellet injection 4.1.1 Shattering 4.1.2 Ablation 4.1.3 Material deposition 4.1.4 Ablation 4.1.5 Material deposition 4.1.6 Ablation 4.1.7 Ablation 4.1.8 Material deposition 4.1.9 Particle and energy balance 4.1 Kinetic equation 4.4 Runaway generation 4.4.1 Kinetic equation 4.4.2 Fluid runaway sources 4.5 Numerical implementation and settings 5.1 First stage deuterium injections 5.1 First stage deuterium injections 5.2 Second stage neon injections 5.2.1 Single vs two-stage injection dynamics 5.2.2 Radiative vs transport losses 5.3.1 Current quench aimulations 5.3.1 Current quench aimulations 5.3.2 Qualitative interpretation based on energy balance 5.3.2 Qualitative interpretation based on energy balance</td>	3.5.2 Pellet material deposition Disruption model 4.1 Shattered pellet injection 4.1.1 Shattering 4.1.2 Ablation 4.1.3 Material deposition 4.1.4 Ablation 4.1.5 Material deposition 4.1.6 Ablation 4.1.7 Ablation 4.1.8 Material deposition 4.1.9 Particle and energy balance 4.1 Kinetic equation 4.4 Runaway generation 4.4.1 Kinetic equation 4.4.2 Fluid runaway sources 4.5 Numerical implementation and settings 5.1 First stage deuterium injections 5.1 First stage deuterium injections 5.2 Second stage neon injections 5.2.1 Single vs two-stage injection dynamics 5.2.2 Radiative vs transport losses 5.3.1 Current quench aimulations 5.3.1 Current quench aimulations 5.3.2 Qualitative interpretation based on energy balance 5.3.2 Qualitative interpretation based on energy balance

1

Introduction

The steadily increasing energy consumption around the world will make future generations demand for reliable and clean energy sources larger than ever before. If successful, fusion energy has the potential to play a key role in a future carbon-free energy system. Even if commercial fusion does not become available soon enough to tackle the most urgent climate changes, it may still be of importance to meet the increasing energy demand during the later half of this century [1].

A fusion reactor will work by fusing two light nuclei, whose combined mass is larger than the resulting fused nuclei [2]. The mass difference is released as energy. The reaction usually consists of a deuterium nucleus (²H) fusing with a tritium (³H) nucleus, forming an α -particle (⁴He) and a neutron. The comparatively light neutron leaves the confined fuel with 80% of the released energy. This energy is then converted to heat that is used to drive a turbine, which generates electricity.

In many ways, fusion can be regarded as an ideal energy source [3]: The fusion reactions do not produce any other waste products than helium, making fusion sustainable from a climate perspective. The hydrogen isotopes used as fuel can be extracted from ordinary sea water, making the fuel reserves practically unlimited. The energy production is not weather-dependent, and the fusion energy output can therefore be adapted to varying demand to a larger extent than weather-dependent renewable energy sources. The half-life of the radioactive material produced by neutron bombardment of the reactor wall is of the order of hundreds of years. This is about a thousand times shorter than the hundreds of thousands of years for some of the high level waste produced by a fission reactor. Finally, fusion produces no by-products that may be used for nuclear weapons.

There are, however, many difficulties involved in realizing a fusion power plant [4]. In order to fuse two nuclei, the distance between them must be similar to their de Broglie wavelength. At this point, the probability to overcome the remaining Coulomb potential barrier by tunneling becomes significant. For this to happen, the nuclei have to overcome the Coulomb repulsion at larger distances, and therefore they have to collide at a very high energy. One way to achieve such energetic collisions is to heat the fuel to a temperature of the order of 10^8 K. Confining such a hot fuel and maintaining the conditions necessary for a substantial fusion reaction rate is a challenging task. The most developed method to overcome these challenges is so-called magnetically confined fusion, which is introduced in section 1.1.

Although magnetically confined fusion devices allow for a stable confinement of the fusion fuel during normal operation, they must be prepared for off-normal events where the confinement is suddenly lost, called disruptions [5]. These events result in a rapid release of the stored thermal and magnetic energy, which can be very be dangerous for the machine integrity. The basic features of disruptions, their consequences, and suggested methods to mitigate their impact on the machine are introduced in section 1.2. However, disruption mitigation in large future fusion reactors still contains a number of open questions. This is the topic to which this thesis aims to contribute, by providing specific guidelines to mitigate the effects of such events.

1.1 Magnetically confined fusion

At the high temperatures necessary for fusion energy generation, the electrons in the fuel are separated from the atomic nuclei. When this happens, the fuel becomes a plasma, which may be regarded as a gas consisting of charged particles. Since the particles are charged, they may be controlled by a magnetic field, making magnetically confined fusion possible.

When a charged particle is subject to a magnetic field, it follows a gyro-motion around the magnetic field lines [2]. Naively, one could therefore confine the plasma by bending the magnetic field into a torus. This is indeed the basic principle behind magnetically confined fusion. However, the construction of such a magnetic field inevitably introduces a gradient of the field strength along the major radius of the torus. For reasons to be described later in section 2.2.2, this magnetic field gradient makes a plasma in a purely toroidal magnetic field intrinsically unstable. To avoid this instability, a poloidal twist of the magnetic field is introduced. The most studied design for such a configuration is the so-called tokamak, where the poloidal twist is achieved by inducing a high current through the plasma, of the order of mega-amperes (MAs).

The tokamak design is utilised by the JET (Joint European Torus) device, currently holding the record for the ratio of the generated fusion power to the supplied heating power - the so-called Q-factor - at 0.67. It is also used for future devices aiming at producing a net gain of fusion energy. These devices include SPARC, under development by Commonwealth Fusion Systems and Massachusetts Institute of Technology [6], STEP (Spherical Tokamak for Energy Production) under development at the Culham Centre for Fusion Energy [7], and ITER, which is currently under construction in France [8]. ITER will be the largest among the next generation of fusion devices, and is the main concern of this thesis. The ITER project is an international collaboration between the European Union, United States, China, Russia, South Korea, India and Japan, aiming for a Q-factor of 10. In order to increase the Q-factor, next-generation devices will store significantly larger thermal and magnetic energies than current experiments. This increases the severity of the threat posed by disruptions, to which we now turn our attention.

1.2 Tokamak disruptions

Disruptions are a type of operational failure, where the plasma confinement is suddenly lost, and the energy contained in the plasma is dissipated to the surrounding structures. The thermal energy may be released over a time scale as short as 0.1-1

ms. In the nuclear operation phase of ITER, this energy will be around 300 MJ. The dissipation time scale for the magnetic energy associated with the plasma current is 1–2 orders of magnitude longer, depending on the size and operation parameters of the tokamak. For ITER-like parameters, the magnetic energy is somewhat larger than the thermal energy, about 1 GJ.

The rapid release of these energies might cause severe damage to the device [5, 9]. The rapid release of the thermal energy can cause a substantial sputtering, or even melting, of the wall material. A rapid current drop can induce large eddy currents in the structures surrounding the plasma. The interaction between these currents and the strong magnetic fields used to control the plasma may result in large forces being exerted on the machine. Conversely, if the current decay is slow, the plasma control might be lost while there is a substantial current remaining in the plasma. The currents in the outer layers of the plasma may then begin to flow through the surrounding structures, forming a so called halo-current. Such a current may also lead to large forces being exerted on the machine. Finally, the potentially most severe threat to a reactor-scale tokamak, is the generation of extremely energetic electrons, called runaway electrons. These electrons may cause a significant melting of wall material upon wall impact. Their existence is made possible by the nature of the energy dependence of the drag force felt by a charged particle moving through the plasma, together with the induction of a large electric field in the plasma during the disruption. These phenomena are described in the next subsection.

1.2.1 Runaway electrons

In a plasma, the collisional interaction between the particles is dominated by the Coulomb force resulting from their charge [4]. This leads to the somewhat counterintuitive phenomenon that the momentum change during a collision decreases with the relative initial momentum of the colliding particles. This may be explained by the fact that particles with a higher relative momentum spend less time in the vicinity of one another, giving the Coulomb force less time to cause a momentum change. As a result, the drag force felt by a charged particle moving through the plasma has a decreasing trend at high momenta.

Consider now the case where the plasma is subject to a strong electric field. The decrease in the drag force at high momenta enables the force exerted by the electric field to overcome the drag force for particles with a momentum higher than a critical value. These particles can therefore be accelerated to extremely high energies, becoming so-called runaways [10]. Typically, the light electrons become true runaways, and are the runaway species we are concerned with here, as ion runaway acceleration is limited by friction with the bulk electrons [11].

During a tokamak disruption, the fast temperature drop leads to the induction of a large electric field, which can lead to the generation of runaway electrons [5]. This electric field induction is due to the relation between the temperature and conductivity of the plasma. The conductivity is inversely proportional to the drag force resulting from the thermal motion of the plasma particles, and therefore the conductivity increases with temperature. Thus, a rapid temperature drop is accompanied by a rapid drop in the conductivity. However, the inductive properties of the plasma current set a lower limit on the time scale over which the current can change. As a result, an electric field must be induced during the fast temperature drop to compensate for the associated drop in the conductivity.

The presence of an electric field only enables the existence of runaway electrons. Actual generation of runaway electrons requires some mechanism that can accelerate electrons above the critical momentum, where the force from the electric field overcomes the drag force. During a tokamak disruption a number of such mechanisms are active.

One runaway generation mechanism which is always active, as long as the electric field is high enough to enable the existence of runaways, is the so-called Dreicer mechanism [12]. It builds on the fact that the distribution of the electron momenta tends to equilibrate towards a Maxwellian. The high-energy tail of the Maxwellian might have particles with momenta higher than the critical one, allowing those electrons to become runaways. When the high-energy tail is accelerated, the rest of the electrons re-equilibrate towards a Maxwellian, so the tail is re-populated. This process repeats itself, allowing a steady runaway generation.

Another runaway generation mechanism builds on the comparatively slow cooling of the hot tail of the pre-disruption momentum distribution [13–15]. Due to their lower collisionality, these hot-tail electrons equilibrate much slower to the much lower post-disruption temperature than the less energetic electrons. As the induced electric field increases, the hot tail of the initial distribution can be caught above the current critical energy for runaway acceleration, before they have time to thermalise. This mechanism, called the hot-tail mechanism, can produce a burst of runaway electrons in the early stages of the disruption.

The nuclear activity in a deuterium-tritium plasma provides two additional mechanisms for runaway generation [16]. The tritium itself undergoes β^- -decay, and the emitted electron might have a momentum above the runaway threshold. Moreover, γ -photons emitted from the activated wall can accelerate electrons above the runaway threshold by Compton scattering.

Finally, an existing runaway population can amplify itself through the so-called avalanche mechanism [17]. An existing runaway electron can collide with a slower electron in such a way that both electrons have a momentum above the runaway threshold after the collision. This mechanism gives rise to an initially exponential growth of the runaway population, and therefore even a very small seed of runaways generated by any of the above mechanisms can be amplified to problematic proportions. This mechanism is exponentially sensitive to the plasma current [18], and is therefore particularly problematic for large, high-current, reactor-scale tokamaks such as ITER.

1.2.2 Mitigation techniques

The potential damage resulting from a disruption described above poses strict requirements for disruption mitigation [5]. Three main tasks must be accomplished by a tokamak disruption mitigation system. The first is to minimize the localized heat loads on the plasma facing components, by spreading the heat loads over as large an area as possible. Secondly, the disruption mitigation system must control the time scale for the current decay so that it is long enough to avoid excessive eddy currents, but short enough to avoid excessive halo currents. Finally, the current carried by runaway electrons impacting the wall should be minimized.

The currently envisaged disruption mitigation method is to inject dissipative material into the plasma as soon as the emerging disruption is detected. The injection primarily consists of hydrogen isotopes and/or noble gases such as neon or argon. As this material enters the plasma in the form of neutral atoms, it may emit atomic line radiation when excited by exposure to the plasma. This radiation can release the thermal energy isotropically, reducing the maximum localised heat loads. The cooling properties of the injected material can also be tuned to gain some control over the temperature after the initial drop, which in turn determines the current decay rate.

Moreover, as the injected material ionizes, the electron density increases. This leads to an increase in the drag force felt by the electrons in the plasma, which reduces the runaway generation. Recent work has however indicated that this suppression of the runaway generation may only work to a limited extent [19]. The reason is that a larger injection is accompanied by a stronger radiative cooling. For sufficiently large injections, the plasma becomes cold enough for the injected ions to recombine with the free electrons, before the runaway generation phase is over. The bound electrons contribute to the number of target electrons for the avalanche mechanism to practically the same extent as free electrons, an effect that is only partially compensated by their contribution to the drag force [20]. For this reason, recombination can substantially enhance the avalanche, which might limit the prospects for runaway suppression by massive material injection. Runaway suppression in high-current tokamaks thus remains an open question.

Conventionally, the injected material is delivered as a gas puff from a pressurized vault [5]. While this technique is comparatively simple, it comes with a number of disadvantages. The injected gas ionizes rapidly when exposed to the still hot plasma. When ionized, the injected material becomes tied to the magnetic field, which substantially slows down the transport towards the plasma core. Moreover, the gas injection introduces a perturbation to the magnetic field, which accelerates the growth of the plasma instabilities whose detection indicates the disruption onset. As a result, the disruption might fully begin before the injected material has reached all parts of the plasma.

Another approach, that can provide better core penetration, is to inject material in the form of a solid, cryogenic pellet. The exposure to the plasma causes the pellet to ablate and deposit its content along its trajectory. If desired, the ablation of the pellet can be made more efficient by shattering the pellet into smaller shards before it enters the plasma, forming a so-called shattered pellet injection (SPI). This technique has been chosen as the baseline for the disruption mitigation system at ITER [21].

The details of the design and the operation parameters for the disruption mitigation system at ITER, as well as other reactor-scale devices, are still an open question. A major difficulty associated with this question is that the solutions to the three main issues to be addressed by the disruption mitigation system to some extent conflict with each other. For example, the radiative dissipation of the thermal energy practically favours as large an injection as possible. On the other hand, a large injection might result in too low a temperature after the temperature drop, leading to a fast current quench time. In addition, rapid cooling increases the hot-tail runaway generation [22, 23].

A recently suggested method attempting to circumvent these issues is to divide the injection into two stages [24]. The first injection would consist of a large amount of pure deuterium. The purpose of this stage would be to cool the plasma through dilution, by a factor 10–100, without significantly perturbing the magnetic field configuration or radiating away any substantial amount of thermal energy. A few milliseconds later, a smaller neon injection would follow, which radiatively dissipates the thermal energy. The division of the cooling into two steps gives the hot tail time to equilibrate at an intermediate temperature before the runaway generation is initiated. In that way, this scheme could potentially lead to an efficient suppression of the hot-tail runaway generation. In addition, if the plasma perturbation destroying the magnetic field configuration is initiated at a lower temperature, the slower thermal motion reduces the thermal energy transport. Instead, a larger fraction of the thermal energy can be dissipated through radiation, reducing the danger of localised wall hot spots.

The results presented in ref. [24] indicate that it is indeed possible to substantially cool the plasma through dilution by a deuterium shattered pellet injection, without destroying the magnetic field configuration. The outcome and optimisation of this two-stage injection scheme, particularly the runaway generation, has however not yet been studied in detail. This task is the main subject of this thesis, and is treated by numerical simulations using the newly developed DREAM (Disruption Runaway Electron Avoidance Model) code [25]. This tool allows for efficient calculations of the time evolution of the spatial profiles of the plasma properties such as the density, temperature and current, as well as the momentum distribution of electrons, during a tokamak disruption. In this project, it has also been equipped with the capability to model disruption mitigation by SPI.

1.3 Thesis outline

The remainder of this thesis is structured as follows: In chapter 2, the underlying plasma physics relevant for the understanding of this work is reviewed. Chapter 3 gives a more detailed description of the physics behind tokamak disruptions and strategies to mitigate them. The disruption model used in this work is described in chapter 4. In chapter 5, this model is used to study disruption mitigation in ITER using a two-stage SPI scheme. As a first step, this study addresses the problem of optimising the pellet shattering for a given pellet size, in order to achieve core penetration while minimizing the amount of material passing through the plasma without being ablated. Using this optimised shattering, the next step is to evaluate the performance of the disruption mitigation system for a range of pellet sizes, in terms of the radiative dissipation of thermal energy, current decay time, and runaway generation. The conclusions are summarised in chapter 6, and areas for further refinement of the model are discussed.

2

Review of the underlying plasma physics

As the fuel in a fusion device is in the plasma state, an understanding of plasma physics is of utmost importance for the study of fusion devices. This chapter therefore gives an overview of the plasma physics necessary for the analysis of tokamak disruptions presented later in this work. In section 2.1, we give a formal definition of a plasma, and discuss some of its basic properties. Section 2.2 describes the motion of charged particles in an electromagnetic field, and how the properties of this motion can be used to confine a plasma in a tokamak. Next, section 2.3 describes the collisional interaction between particles in a plasma, and scaling laws for the collisional features relevant during tokamak disruptions are derived. Finally, section 2.4 introduces the theoretical framework underlying the detailed plasma modelling performed in this work.

2.1 Definition and basic properties

If a gas is heated to a high enough temperature, the electrons will eventually be separated from their atomic nuclei. The resulting state is called a plasma, and is often referred to as the fourth state of matter, following the solid, liquid and gas state [4]. In order to achieve deuterium-tritium fusion, the fuel must be heated to about $1.5 \cdot 10^8$ K, which is well above the temperature needed to enter the plasma state [2]. In plasma physics, the temperature is usually measured in electronvolts (eV) referring to the typical energy $k_{\rm B}T$ of the thermal motion of the particles, where $k_{\rm B}$ is the Boltzmann constant. A temperature in Kelvin differs by a factor of $e/k_{\rm B} = 11592$ K/eV from the temperature given in eV, where e is the elementary charge.

In order to attain the characteristic properties of a plasma, however, the set of particles must have the properties called *quasi-neutrality* and *collective behavior* [4]. These conditions set constraints on the length scale, time scale and particle density. Collective behaviour refers to the fact that charged particles in a plasma interact with many other particles at the same time through long-range electromagnetic forces. This situation is very different from neutral particles, which mostly interact through close, pair-wise, collisions. To fulfill the definition of a plasma, the long-range interactions must dominate over the close collisions with neutral particles. For this to be the case, the collision frequency with neutral particles must be low compared to the frequency of the fundamental oscillations in the plasma.

Quasi-neutrality means that the net charge of the plasma is neutral over macro-

scopic length scales, although charge imbalances occur at microscopic scales. In a plasma, this is achieved as the particles surrounding an emerged charge imbalance react to this imbalance due to the Coulomb force, rearranging themselves in a way that screens out the imbalance as seen from some distance away. This phenomenon is called *Debye screening* [4]. If one, for instance, introduces a positive point charge in a plasma, it will attract the nearby electrons, and in that way increase the electron density in the vicinity of the point charge. The negative charge of the extra electrons will then cancel the positive charge of the point charge except at very short distances. The length scale for this screening can be calculated assuming the electrons follow a Boltzmann distribution. In that case, their density is proportional to $\exp(e\phi/(k_{\rm B}T))$ where ϕ is the electric potential. The electric potential as a function of the distance r from the point charge can be shown to follow the Coulomb potential multiplied by a factor of $\exp(-r/\lambda_{\rm D})$, where [4]

$$\lambda_{\rm D} = \sqrt{\frac{\epsilon_0 k_{\rm B} T}{n_0 e^2}}.$$
(2.1)

Here, ϵ_0 is the dielectric constant and n_0 is the background (unperturbed) ion particle density. The quantity λ_D is called the Debye length, which is a typical length scale for the screening of the point charge. In order for quasi-neutrality to occur, the length scale of the plasma must be much longer than this length scale. Moreover, for the Boltzmann statistics to be valid, the number of particles within a distance λ_D must be large, which poses a constraint on the particle density.

2.2 Magnetically confined plasmas

Due to the high temperatures in a plasma, they can not be confined in ordinary containers. However, since the particles in a plasma are charged, it is possible to control them by the use of magnetic fields. The goal of this section is to provide an understanding of how this is done in devices called tokamaks, which are the most common design for present day fusion experiments. Although the collective behaviour of the particles in a plasma is essential for the understanding of magnetic confinement, it is instructive to start with understanding the motion of a single particle in prescribed electromagnetic fields. This is the subject of section 2.2.1. We then proceed by applying these phenomena to the circumstances representative of a tokamak in section 2.2.2.

2.2.1 Charged particle motions in electromagnetic fields

To start with, we consider the motion of a charged particle in a homogeneous and constant magnetic field, without the presence of an electric field. A particle of charge q moving with velocity v in a magnetic field B is subject to the Lorentz force F = qvB, directed perpendicularly to both the magnetic field and the motion of the particle [4]. In the absence of other forces, the particle will then follow a circular orbit in the plane perpendicular to the magnetic field, with a centripetal force equal to the Lorentz force, while moving freely along the magnetic field lines. This tendency of charged particles to stick to a fixed circular orbit around the magnetic field lines forms the foundation of plasma confinement by the use of magnetic fields [2]. The circular motion is called *gyro motion* or *Larmor motion*, and can be shown to have a rotation angular frequency of [4]

$$\omega_L = \frac{|q|B}{m} \tag{2.2}$$

and a radius [4, p. 20]

$$r_L = \frac{mv_\perp}{|q|B},\tag{2.3}$$

where v_{\perp} is the velocity perpendicular to the magnetic field.

In more general electromagnetic field settings, however, the center of the gyro motion, called the *guiding center*, might be subject to drifts across the field lines. We now consider the case where we add a homogeneous, constant, non-zero electric field. The component of the electric field in the plane perpendicular to the magnetic field will then affect the gyro motion differently during different parts of the gyro motion. When the particle moves in the same direction as the electric field it decreases. Since the change in speed is accompanied with a proportional change in the radius of the gyro motion, the radius of the gyro motion will be larger during one half of the gyrations than the other. Averaged over a whole gyration, this causes a net drift called the $E \times B$ drift (pronounced "E cross B drift") of the guiding center perpendicular to both the magnetic and electric field. This process is illustrated in the left panel of figure 2.1. Quantitatively, the guiding center velocity of this drift can be shown to be [4]

$$\mathbf{v}_{E\times B} = \frac{\mathbf{E}\times\mathbf{B}}{B^2}.\tag{2.4}$$

An important feature of this drift velocity is that it is independent of the particle charge, because the charge dependence of the Lorentz force and the force from the electric field cancel each other.

A similar drift emerges if there is a gradient in the magnetic field strength, with the important difference that this drift is charge-dependent. A gradient in the magnetic field makes the radius of the gyro motion longer in the part of the gyration where the magnetic field is weaker compared to the part where it is stronger. This results in a drift perpendicular to both the magnetic field and its gradient, as illustrated in the right panel of figure 2.1. This drift is referred to as the ∇B -drift (pronounced "grad B-drift"). Quantitatively, this drift can be shown to be given by [4]

$$\mathbf{v}_{\nabla B} = \frac{m v_{\perp}^2}{2Bq} \frac{\mathbf{B} \times \nabla B}{B^2}.$$
 (2.5)

2.2.2 The tokamak design

Based on the principle that charged particles follow magnetic field lines, an intuitively appealing way to confine a plasma is to establish a magnetic field with closed field lines, such as a toroidal field. This is also the basic principle behind magnetic



Figure 2.1: Particle trajectories for an ion and electron, starting from the origin, with different prescribed electric and magnetic field settings: (Left) a constant uniform magnetic field pointing in the positive z-direction (out from the paper) and a constant uniform electric field pointing in the positive y-direction, giving rise to an $E \times B$ -drift, and (right) a constant magnetic field pointing in the positive z-direction with a uniform gradient in the negative y-direction, giving rise to a ∇B -drift. The gyro-radii and speeds are chosen equal for visibility.

confinement fusion devices, including the so-called *tokamak* [2]. However, due to the particle drifts explained in the previous subsection, a purely toroidal magnetic field is not sufficient to confine a plasma.

When bending a magnetic field with straight field lines into a torus, the field line density becomes larger in the inner part of the torus, referred to as the *high field side*, compared to the outer part, referred to as the *low field side*. The drifts introduced by this gradient make ions and electrons drift vertically in opposite directions. The separation of positive and negative charges gives rise to a vertical electric field, which in turn gives rise to an $E \times B$ -drift. This $E \times B$ -drift is charge independent, and makes both ions and electrons drift towards the low field side, out of confinement. This mechanism is illustrated in figure 2.2.

This issue can be resolved by introducing a poloidal twist to the magnetic field, i.e. a magnetic field component circulating the short way around the plasma core, as illustrated in the left panel of figure 2.3 [2]. Such a twist makes the particles circulate around the plasma in the poloidal direction in a way that averages out the effect of the vertical ∇B -drift; in the upper part of the poloidal motion, the ∇B -drift drives the particles away from the plasma center, while in the lower part the ∇B -drift drives the particle back towards the center. Such a trajectory is illustrated in the right panel of figure 2.3. In that way, the charge separation, and hence the $E \times B$ -drift that breaks the confinement, is avoided.



Figure 2.2: Illustration of the drifts causing the loss of confinement in a purely toroidal magnetic field configuration. (Left) The bending of the magnetic field gives rise to a gradient towards the inner part of the torus, causing a ∇B -drift, making positive charges accumulate in the upper part and negative charges in the lower part of the plasma. (Right) The charge separation gives rise to a vertical electric field; the corresponding $E \times B$ -drift makes the plasma drift out of the confinement.

In a tokamak, this twist of the field lines is achieved by driving a large toroidal current through the plasma. In such a configuration, the magnetic field lines can be shown to circulate in the plasma following nested toroidal-like surfaces called *flux surfaces*, sketched in the left panel of figure 2.3 [2]. The term flux surfaces reflects the fact that these surfaces are surfaces of constant poloidal magnetic flux (and also approximately constant pressure), so that they can be uniquely characterised by this flux.

The current is most commonly driven by a transformer-like mechanism, where the tokamak plays the role of the secondary coil; a time-varying vertical magnetic field is generated through the tokamak, which drives a current through the plasma by induction [2]. In order to drive a DC current in the tokamak, however, the current in the primary coil generating the vertical magnetic field has to vary linearly, which can only be done for a limited amount of time. The operation of a tokamak is therefore necessarily divided into several sessions, referred to as *pulses*, *shots* or *discharges*. Compared to everyday-values, the plasma currents in a tokamak are very large. In the largest present day tokamak called JET (Joint European Torus), the plasma current is approximately 2 MA, and ITER plans for an operation at a plasma current of 15 MA [8].

Such large plasma currents can cause serious problems, due to the large energies associated with them [26]. In the case of a disruption, a form of operational failure where the plasma suddenly cools, these energies can cause severe damage to the surrounding structures. The mitigation of the impact of such events is the main topic of this work, and a detailed background about disruptions is given in chapter 3. In order to understand the phenomena at play during disruptions, it is however



Figure 2.3: (Left) Illustration of the magnetic field configuration in a tokamak; the magnetic field is the sum of a toroidal component generated by external coils (which are not shown in the figure) and a poloidal component generated by the plasma current. The resulting magnetic field lines follow helical paths around the plasma, ordered in a set of nested flux surfaces. (Right) Projection of the guiding center trajectory of an ion in a tokamak on the poloidal plane, illustrating how the ∇B -drift is averaged out by the motion along the poloidal component of the magnetic field. A flux surface with circular cross section, tangential to the ion guiding center trajectory, is included as reference. The drift away from the flux surface is exaggerated for visibility.

important to understand another branch of plasma physics, namely the physics of collisions in a plasma, which is therefore the topic of the next section.

2.3 Collisions in a plasma

As we will see, collisions in a plasma differ quite significantly from collisions in a gas. Instead of the pair-wise close collisions dominating in a gas, the collisional dynamics in a plasma is dominated by long-range interactions through the Coulomb force. An important consequence of this is that, as we will see shortly, fast (or hot) particles are *less* collisional than slow (or cold) particles, as opposed to the case in a gas. In this section, the physics of Coulomb collisions is illustrated, and related to the electrical conductivity of the plasma and eventually to the runaway electron phenomenon.

2.3.1 Coulomb collisions

A qualitative comparison of the difference between Coulomb collisions and close collisions prevailing in a gas is illustrated in the left panel of figure 2.4. This illustration is made using a simple simulation of a particle traveling through a fixed ion background (corresponding to background particles with a much larger mass than



Figure 2.4: (Left) Illustration of a neutral atom (black) and two electrons (blue) with different velocities moving through a background of fixed ions, comparing electron-ion Coulomb collisions with close collisions with neutrals, and illustrating the velocity dependence of Coulomb collisions. (Right) Illustration of a single Coulomb collision between an incident electron and a fixed ion.

the test particle, which would be the case if the moving particle was an electron). If the particle is neutral, so that the interaction with the background takes place through close collisions, the resulting trajectory will be made up of straight line segments with sharp deflections where a collision takes place. If the particle is an electron, on the other hand, the long-range Coulomb interaction with the background ions cause the electron to follow a smoothly curved path. For the neutral particle, the deflection angle is independent of the particle speed, but for the electron interacting through the Coulomb force, we see that a fast electron is significantly less affected by the background ions than a slower one. This happens because the time spent in the vicinity of the ions decreases with an increasing electron speed, without affecting the interaction force.

The setup for a single Coulomb collision for an electron colliding with a fixed ion is shown in the right panel of figure 2.4. The ion is assumed here to have an equal but opposite charge compared to the electron. In the absence of the Coulomb force, the electron would have a distance of closest approach equal to b as defined in the figure, which is referred to as the *impact parameter*. After the collision, the electron is deflected by an angle α compared to its original direction.

The velocity scaling of the cross section and collision frequency of electron-ion Coulomb collisions, defined as the rate at which the electron momentum undergoes an order unity change, can be determined as follows. For Coulomb collisions, the momentum change is dominated by the accumulated effect of many small-angle interactions [10]. However, as we are primarily interested in the velocity scaling here, we may simplify the picture by considering events where the order unity momentum change occurs due to a single Coulomb interaction. During the main part of this interaction, the Coulomb force is roughly given by

$$F_{\rm c} \sim \frac{e^2}{4\pi\epsilon_0 b^2},\tag{2.6}$$

and is felt during a collision duration time roughly given by $t_{\rm c} \sim b/v$. From these

considerations, we may relate the impact parameter necessary for giving a significant deflection of the electron to the electron velocity according to

$$\Delta p \sim mv \sim F_{\rm c} t_{\rm c} \sim \frac{e^2}{4\pi\epsilon_0 bv} \Rightarrow b \sim \frac{e^2}{4\pi\epsilon_0 mv^2}.$$
(2.7)

The cross section can then be estimated according to

$$\sigma_{\rm c} \sim \pi b^2 \sim \frac{e^4}{16\pi\epsilon_0^2 m^2 v^4}.$$
 (2.8)

The collision frequency can now be estimated by the number of electrons passing through a cross section of area σ_c per unit time, according to

$$\nu \sim nv\sigma_{\rm c} \sim \frac{ne^4}{16\pi\epsilon_0^2 m^2 v^3}.\tag{2.9}$$

Here, n is the density of electrons or ions, assumed to be equal here for simplicity. The important take-away here is the $1/v^3$ -scaling, which quantifies the previous observation that fast particles are less collisional in a plasma. In practice, however, the particles in a plasma do not all have the same velocity, but usually a thermal, Maxwellian, distribution of velocities. To estimate the collision frequency for a thermal plasma, we may therefore replace the velocity v by the thermal speed $v_{\rm th} = \sqrt{k_{\rm B}T/m}$ and obtain the thermal collision frequency as

$$\nu_{\rm th} = \frac{ne^4}{16\pi\epsilon_0^2 m^2 v_{\rm th}^3} = \frac{ne^4}{16\pi\epsilon_0^2 m^{1/2} (k_{\rm B}T)^{3/2}},\tag{2.10}$$

where the main takeaway is that the thermal collision frequency scales as $T^{-3/2}$.

2.3.2 Conductivity

Having estimated the collision frequency, we may determine a scaling law for the plasma conductivity. If an external electric field E is applied to the plasma, the Maxwellian distribution will be slightly shifted, so that it has a non-zero average velocity δv . This shift in the average velocity is related to the current density according to $j = ne\delta v$. The average velocity may be estimated by equating the accelerating force per electron due to the electric field with the collisional drag force. Recalling that the collision frequency is defined as the rate at which an electron undergoes an order unity momentum change, we obtain the scaling $F_{\rm drag} \sim m \delta v \nu$ for the drag force. As collisions among electrons do not alter the total electron momentum (they only transfer momentum from one electron to another), electronelectron collisions do not affect the scaling law of interest here (although they do affect the details of the electron velocity distribution, and should therefore be included in a complete analysis [10]). For small $\delta v/v_{\rm th}$, the relative velocity compared with the ions is dominated by the thermal motion, and we therefore use the thermal collision frequency in the expression for the drag force. With these considerations, we may estimate δv according to

$$eE = F_{\rm drag} \sim m \delta v \nu_{\rm th} \Rightarrow \delta v = \frac{eE}{m\nu_{\rm th}},$$
 (2.11)

which gives a relation between the electric field and the current according to

$$j = ne\delta v \sim \frac{ne^2}{m\nu_{\rm th}} E = \sigma E.$$
(2.12)

Here, we have introduced the conductivity σ , which is given by

$$\sigma \sim \frac{ne^2}{m\nu_{\rm th}} \sim \frac{16\pi\epsilon_0^2 (k_{\rm B}T)^{3/2}}{m^{1/2}e^2}.$$
 (2.13)

An interesting note here is that the conductivity is independent of the density, meaning that it is independent of the number of charge carriers, which might seem surprising. This can be understood as even though the number of charge carriers increases with the density, the drag force also increases with the density, and these effects cancel each other.

A more complete analysis confirms the most important features of equation (2.13), such as the density independence and the $T^{3/2}$ -scaling [10]. There are however a few details missing in this scaling law. If multiple ion species are present, a factor $1/Z_{\text{eff}}$ is introduced, where the *effective charge number* Z_{eff} is defined as [10]

$$Z_{\text{eff}} \equiv \frac{\sum_{i} \sum_{j} n_{ij} Z_{ij}^2}{\sum_{i} \sum_{j} n_{ij} Z_{ij}}.$$
(2.14)

The effective charge number is a measure of the average charge number (charge given in terms of the elementary charge) of all the background ions, n_{ij} and Z_{ij} are the density and charge number respectively of charge state i of ion species j. Moreover, if the effect of electron-electron collisions is taken into account, another dimensionless factor of order unity is introduced, which is equal to 0.51 for $Z_{\text{eff}} = 1$ and approaches 1 for large Z_{eff} [10]. Finally, the conductivity is affected by the fact that the momentum change is dominated by the accumulated effects of many small-angle collisions, rather than the single large-angle collisions considered when obtaining the scaling law. While this does not alter the scaling with the dimensional properties, it introduces a dimensionless factor $\ln \Lambda$, called the *Coulomb logarithm*, to the collision frequency, leading to a factor $1/\ln \Lambda$ in the conductivity [4]. Here, Λ is defined as $\Lambda = \langle \lambda_{\rm D}/b \rangle$, that is, the average over the Maxwellian velocity distribution of the ratio $\lambda_{\rm D}/b$ with the impact parameter b related to the velocity according to equation (2.7). The Debye length enters here as an estimate of the largest possible impact parameter before the incident particle will no longer feel the charge of the target particle due to screening. As the deflection angle decreases with an increasing impact parameter, the Coulomb logarithm provides a measure of the impact of small angle collisions compared to large angle collisions. Although Λ depends on the temperature and density, the logarithm is rather insensitive to the plasma parameters, and the Coulomb logarithm usually remains in the range 10-20 for a very wide range of plasma parameters [4].

The temperature scaling of the conductivity plays an important role during a disruption, as it means that a drop in temperature is accompanied by a drop in the conductivity. The conductivity after the temperature drop has a major impact on the current dissipation rate. The current dissipation rate has in turn a significant

impact on the electromagnetic forces exerted on the machine, as described later in section 3.2. In addition, as the current in an inductive system can not change instantly, a fast drop in the conductivity has to be compensated by a large induced electric field. This is quite similar to the high voltages that may occur when one disconnects the power supply to an ordinary coil, as the reader might be familiar with from circuit theory. The electric field thus generated can quickly accelerate the light electrons in the plasma to extremely high energies, forming a beam of so-called runaway electrons, as described next.

2.3.3 The runaway phenomenon

We now turn our attention to the drag force felt by a single electron traveling with velocity v through a plasma consisting of both electrons and ions. For $v \ll v_{\rm th}$, the relative velocity, and hence the collision frequency, as seen by the particle traveling through the plasma is dominated by the thermal motion of the plasma particles. In such a case, we expect a linear dependence of the drag force on v, similarly to the case with the shifted velocity distribution discussed in section 2.3.2. On the other hand, for $v \gg v_{\rm th}$, the relative velocity is dominated by the velocity v, so that $\nu \propto 1/v^3$. As the drag force scales as $F_{\rm drag} \sim mv\nu$, we now expect a $1/v^2$ -scaling of the drag force. A more complete analysis, outlined in section 2.4.1.1, confirms these asymptotic scalings. More precisely, it can be shown that the dependence of the drag force on the velocity for a non-relativistic electron is proportional to the Chandrasekhar function,

$$G(x) = \frac{\operatorname{erf}(x) - x \operatorname{erf}'(x)}{2x^2},$$
(2.15)

where erf is the error function and $x = v/v_{\text{th}}$ [10]. This function is illustrated in figure 2.5.

It is instructive here to emphasise the contrast between the drag force in a plasma and the drag force in an ordinary fluid. In an ordinary fluid, the drag force increases monotonically with the velocity, typically linearly or quadratically. In such a case, if an accelerating force, such as gravity, is present, an object moving through the fluid will accelerate until the drag force equals the accelerating force. At this point, the acceleration will stop, at a finite velocity. This is also the case in a plasma at low velocities. At high velocities, however, the drag force decreases as a function of the velocity. This means that, in the presence of an accelerating electric field E, once a particle has gained a speed higher than a critical speed marked as v_c in figure 2.5, the drag force will never balance out the accelerating force. The particle can then continue accelerating to extremely high energies. This phenomenon is called the *runaway phenomenon*, and electrons with a speed larger than v_c are called *runaway electrons* [10]. For such electrons, the energy gain is only limited by the energy losses due to synchrotron radiation and bremsstrahlung, which become significant at very high energies [27, 28].

If the electric field becomes large enough, the accelerating force corresponding to the electric field will be stronger than the maximum friction force. When this happens, all electrons become runaways. An estimate of the electric field required



Figure 2.5: Illustration of the drag force felt by an electron moving through a plasma as a function of velocity. The presence of an accelerating electric field E creates a runaway region at velocities larger than v_c , where the drag force becomes weaker than the accelerating electric force. The Dreicer field E_D and the critical electric field E_c are also marked.

for this to happen can be obtained by inserting $v = v_{\rm th}$ in the asymptotic expression for the drag force at high velocities. The electric field $E_{\rm D}$ corresponding to this force is called the *Dreicer electric field* and is given by

$$E_{\rm D} = \frac{e^3 n \ln \Lambda}{4\pi\epsilon_0^2 T}.$$
(2.16)

When considering the full expression for the drag force, it can be shown that the electric field at which all electrons become runaways is approximately equal to $0.21E_{\rm D}$ [10].

As the electron speed is limited by the speed of light, one may anticipate from the classical treatment that the drag force does not go all the way down to zero at high velocities. This sets a lower limit on the electric field necessary for the existence of runaway electrons. A relativistic treatment in the high energy limit shows that, in a fully ionized plasma, the drag force approaches the force corresponding to a critical electric field equal to [10],

$$E_{\rm c} = \frac{e^3 n \ln \Lambda}{4\pi\epsilon_0^2 mc^2}.\tag{2.17}$$

During normal tokamak operation, due to the very high conductivity, an electric field of the order of 1 mV/m is sufficient to drive the plasma current, which is usually not high enough to enable runaway generation. The situation might however be different in the case of a disruption, where the plasma suddenly cools [26]. Typically, the temperature drops by about three orders of magnitude, resulting in a decrease in the conductivity by a factor of the order of $10^{-4} - 10^{-5}$ (recall the $\sigma \sim T^{-3/2}$ scaling from section 2.3.2). On the short time scale of the temperature drop, the current density is essentially constant, so that $E \sim 1/\sigma$. The electric field thus increases by

a factor of $10^4 - 10^5$, and can become much larger than E_c , so that runaways are generated.

Once a runaway current is formed, the low collisionality at high energies allows the runaway current to remain in the plasma for a comparatively long time, even after the electric field has decayed. Therefore, it is likely that a significant part of the runaway current will crash into the wall before it is dissipated, potentially causing severe damage to the device [5]. The details of the mechanisms for runaway generation during tokamak disruptions, other potential damage associated with these events, and possible mitigation techniques, are described in the next chapter. However, before narrowing down to the details of disruptions, we conclude this chapter with a section giving an overview of the theoretical framework for plasma modeling. This framework lays the foundation for the model used in this work, which is described in chapter 4.

2.4 Theoretical models of plasmas

The most exact way to model a plasma would be to solve the equations of motion for every particle in the plasma individually. This approach is however unfeasible due to the large number of particles involved ($\sim 10^{23}$ in an ITER-like plasma). The complexity must therefore be reduced to a statistical description, where only the distribution of particles, and not every particle individually, is modelled. This can be done with various levels of sophistication, depending on how many dimensions of the distribution are modelled in detail, and how many are only included through moments of the distribution over that degree of freedom. Two distinct classes of models are *kinetic models*, where both the velocity and configuration space distributions are resolved, and *fluid models*, resolving only the configuration space dynamics. These descriptions are outlined in this section, together with the equations describing the electromagnetic fields that are also a crucial part of plasma modeling.

2.4.1 Kinetic models

In a kinetic model, a particle species in a plasma is described by the distribution $f(\mathbf{r}, \mathbf{v}, t)$ [4]. As \mathbf{r} and \mathbf{v} both depend on three variables each, the distribution function is a function of 7 variables. It is customary to normalise the distribution function so that its integral over the velocity space gives the particle density,

$$n(\mathbf{r},t) = \int f(\mathbf{r},\mathbf{v},t)d\mathbf{v}.$$
 (2.18)

With this normalisation, the number of particles in the infinitesimal phase space volume $dxdydzdv_xdv_ydv_z$ located at $\mathbf{r} = (x, y, z)$ and $\mathbf{v} = (v_x, v_y, v_z)$ at time t is given by

$$dn = f(x, y, z, v_x, v_y, v_z, t) dx dy dz dv_x dv_y dv_z.$$

$$(2.19)$$

The evolution of the distribution function is governed by the equation referred to as the *Boltzmann equation* or *kinetic equation*, which has the form [4]

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f + \frac{\mathbf{F}}{m} \cdot \frac{\partial f}{\partial \mathbf{v}} = \left(\frac{\partial f}{\partial t}\right)_{c} + S.$$
(2.20)

Here, ∇ is the gradient in position space, $\partial/\partial \mathbf{v}$ is the gradient in velocity space, and \mathbf{F} is the force acting on the particles, usually the electromagnetic force $\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$. The right hand side contains the effect on the distribution function caused by collisions via the so-called *collision operator* $\left(\frac{\partial f}{\partial t}\right)_c$. The source term S represents changes in the total number of particles, for example caused by ionization or recombination (where a free electron is caught by an ion). The left hand side may be recognised as the the total time derivative df/dt of the distribution function, i.e. the time derivative in a frame of reference following the particles in the six-dimensional phase space, by invoking the chain rule, and recognising that $d\mathbf{v}/dt = \mathbf{F}/m$ from Newton's second law. The kinetic equation can therefore be regarded as a mathematical formulation of the rather intuitive statement that the distribution function in the particle frame can only change due to collisions and sources. The collision operator is usually rather complicated, and additional details of the form used in this work will be described in chapter 4.

In principle, a complete kinetic description of a plasma would include one distribution function for every species present in the plasma, coupled by their self consistent electromagnetic field and collisional processes. However, despite the very large reduction in complexity compared to following every particle individually, a complete six-dimensional model of a plasma is usually also unfeasible. In this work, we will therefore model the ions, and some aspects of the electron dynamics, by the use of fluid models. Such models are described in section 2.4.2. Before this, we conclude the overview of the kinetic description by introducing the fundamental aspects of the collision operator for Coulomb collisions relevant for the kinetic part of the model used in this work.

2.4.1.1 The Coulomb collision operator

When deriving an expression for the collision operator for Coulomb collisions, one may exploit the fact that the dynamics is dominated by small angle collisions, as discussed in section 2.3.2. This makes it possible to describe the dynamics using a *Fokker-Planck collision operator*, which has the form [10]

$$\left(\frac{\partial f}{\partial t}\right)_{c} = C(f) = \sum_{k} \sum_{l} \frac{\partial}{\partial v_{k}} \left[-\frac{\langle \Delta v_{k} \rangle}{\Delta t} f + \frac{\partial}{\partial v_{l}} \left(\frac{\langle \Delta v_{k} \Delta v_{l} \rangle}{2\Delta t} f \right) \right].$$
(2.21)

The first term can be interpreted as a friction force, and the second term describes a diffusion process in velocity space. The expression $\langle \Delta v_k \rangle$ should be understood as the expectation value of the velocity change along direction $k \in (x, y, z)$ during a short time Δt along a particle trajectory, due to the collisional contribution from all other particles. These expectation values can be calculated based on a detailed calculation of the velocity change in the single particle collision considered in section 2.3.1, assuming small deflection angles. Next, the number of collisions with a given velocity and impact parameter of the incident particles is calculated, in terms of the distribution function of the incident species. Once these quantities are known, the expectation value for the total velocity change can be obtained by integrating their product over velocity space and impact parameters. The resulting contribution to the velocity change of species a due to collisions with species b, for a non-relativistic

plasma, can be shown to be [10]

$$\frac{\langle \Delta v_k \rangle^{ab}}{\Delta t} = -\frac{L^{ab}}{4\pi} \left(1 + \frac{m_a}{m_b} \right) \int \frac{u_k}{u^3} f_b(\mathbf{v}') d^3 v', \qquad (2.22)$$

and

$$\frac{\langle \Delta v_k \Delta v_l \rangle^{ab}}{\Delta t} = -\frac{L^{ab}}{4\pi} \int U_{kl} f_b(\mathbf{v}') d^3 v'.$$
(2.23)

Here, $L^{ab} = (q_a q_b)^2 / (m_a \epsilon_0)^2 \ln \Lambda$, $\mathbf{u} = \mathbf{v} - \mathbf{v}'$ is the relative velocity between the particles of species a and b, $u = |\mathbf{u}|$ and $U_{kl} = (u^2 \delta_{kl} - u_k u_l)/u^3$ where δ_{kl} is the Kronecker delta. If there are multiple species $b_1, b_2, ..., b_n$ colliding with species a, the total collision operator for species a is given by the sum

$$C_a(f_a) = \sum_{k=1}^{n} C_{ab_k}(f_{b_k}).$$
(2.24)

The notation C_{ab_k} should be interpreted as the collision operator for collisions of species a against species b_k , and f_a , f_{b_k} are the distribution functions for species a and b_k , respectively.

In many cases, equation (2.22) and (2.23) can be simplified by certain assumptions for the distribution function of the background species. One such case is collisions between particles of disparate speeds, e.g. collisions between electrons (denoted e) and ions (denoted i) with similar temperatures. The large mass ratio makes the ions move much slower, and the distribution function f_i can therefore be approximated by a delta function around v' = 0. The resulting collision operator for electron-ion collisions takes the form [10]

$$C_{ei} = \nu_{ei}(v)\mathcal{L}(f_e), \qquad (2.25)$$

where the electron-ion collision frequency is $\nu_{ei}(v) = n_i Z_{\text{eff}} e^4 / (4\pi m_e^2 \epsilon_0^2 v^3)$ and $\mathcal{L}(f_e)$ is the Lorentz operator. This operator is proportional to the angular part of the Laplacian operator in velocity space, meaning that C_{ei} describes a diffusion process on a sphere in velocity space at constant speed. As a result, only the direction of the electrons are changed by collisions with the massive ions.

The above approximation based on the disparate speeds of the colliding particles is not valid for electron-electron collisions. However, in many cases, such as the ones considered in this thesis, the electron distribution is dominated by a Maxwellian bulk population, which can be used to simplify the collision operator. Thus, we divide the electron distribution function into a Maxwellian part f_{e0} and a non-Maxwellian part $f_{e1} = f_e - f_{e0}$ containing far fewer particles than f_{e0} . The collision operator can be shown to be bilinear [10], so that the collision operator can be expanded as

$$C_{ee}(f_{e0} + f_{e1}, f_{e0} + f_{e1}) = C_{ee}(f_{e0}, f_{e0}) + C_{ee}(f_{e0}, f_{e1}) + C_{ee}(f_{e1}, f_{e0}) + C_{ee}(f_{e1}, f_{e1}).$$
(2.26)

The first term vanishes, as the collision operator for collisions between two Maxwellians with the same temperature and average velocity is zero [10]. If f_{e1} is small compared to f_{e0} , the last term can also be neglected. The two remaining terms now comprise a *linearised electron-electron collision operator*. The first of these remaining terms,

called the *field particle term*, describes the effect on the Maxwellian bulk from collisions with the small non-Maxwellian population. The second, called the *test particle term*, describes the effect on the non-Maxwellian population from collisions with the Maxwellian. This term can be evaluated by inserting the expression for the Maxwellian population into equation (2.22) and (2.23). The resulting collision operator becomes [10]

$$C_{ee}(f_{e1}, f_{e0}) = \nu_{\rm D}^{ee} \mathcal{L}(f_{e1}) + \frac{1}{v^2} \frac{\partial}{\partial v} \left[v^3 \left(\frac{1}{2} \nu_{\rm s}^{ee} f_{e1} + \frac{1}{2} \nu_{||}^{ee} v \frac{\partial f_{e1}}{\partial v} \right) \right].$$
(2.27)

Here, the terms describe angular scattering, frictional drag, and a parallel velocity space diffusion, respectively. The corresponding collision frequencies are

$$\nu_D^{ee} = \hat{\nu}_{ee} \frac{\operatorname{erf}(x) - G(x)}{x^3}, \qquad (2.28)$$

$$\nu_S^{ee} = 4\hat{\nu}_{ee} \frac{G(x)}{x},\tag{2.29}$$

$$\nu_{||}^{ee} = 2\hat{\nu}_{ee} \frac{G(x)}{x^3},\tag{2.30}$$

$$\hat{\nu}_{ee} = \frac{n_e e^4 \ln \Lambda}{4\pi \epsilon_0^2 m_e^2 v_{\rm th}^3}.$$
(2.31)

We may note here that the average friction force on a particle, $-m_e v \nu_S^{ee}$, is proportional to the Chandrasekhar function (which was illustrated in figure 2.5). This fact gives rise to the runaway phenomenon, as discussed in section 2.3.3.

To obtain a model suitable for the description of the relativistic particles considered in this work, some further extensions of the collision operator must be made. To study the runaway acceleration of particles up to momenta larger than m_{ec} , relativistic effects have to be included. The effects for partially ionized ions and atoms present during a mitigated disruption must also be taken into account. Many of these effects can be cast as extensions to the form of the collision frequencies appearing in equation (2.27). The partially ionized ions and atoms do however not only affect the dynamics of the non-Maxwellian population, but also introduce radiative losses of thermal energy affecting the Maxwellian population. When using a linearised test particle collision operator, as in this work, the dynamics of the Maxwellian population is not accounted for by the kinetic equation. Instead, the Maxwellian population can be described using fluid models, which are introduced next.

2.4.2 Fluid models

In a fluid model, the configuration space dynamics is fully resolved, but the velocity space dynamics are only taken into account through *moments* of the velocity distribution. In general, the moments can be thought of as integrals or averages over the velocity space of various velocity-related quantities, and the *i*th order moment is proportional to the average of \mathbf{v}^i . The zeroth order moment is simply the density (one for each species present in the plasma),

$$n = \int f d\mathbf{v}.$$
 (2.32)

The first order moment is the fluid velocity \mathbf{u} given by

$$\mathbf{u} = \frac{1}{n} \int \mathbf{v} f d\mathbf{v},\tag{2.33}$$

which is related to the species' current density according to $\mathbf{j} = ne\mathbf{u}$. The second order moment is the energy density W given by

$$W = \frac{1}{n} \int \frac{1}{2} m v^2 f d\mathbf{v}, \qquad (2.34)$$

which is related to the species' temperature according to W = 3nT/2. In principle, one could continue expanding the fluid model by taking into account arbitrarily high order moments, but in this work we will restrict ourselves to those above.

The equation governing the evolution of these quantities can be derived by taking the corresponding moment of the kinetic equation. For the density, this yields a continuity equation,

$$\frac{\partial n}{\partial t} + \nabla \cdot (n\mathbf{u}) = S. \tag{2.35}$$

The source term S should now be interpreted as the velocity space integral of the source term in the kinetic equation. In this work, we will however only consider particles following the magnetic field lines in a cylindrically symmetric magnetic field, and therefore the transport term $\nabla \cdot (n\mathbf{u})$ will be neglected. The density evolution of various species is instead governed by the source term, taking into account ionization/recombination processes and external sources. It is also possible to treat sub-domains of the velocity space as fluids, in which case the velocity space volume integrated over when calculating the moments is limited to these sub-domains. This can for example be done for the runaway electrons, which have reached high energies and move at near the speed of light. In such a case, acceleration across the boundary into this region of velocity space gives rise to a density source term. The details of the source terms used in this work will be described in later chapters.

Considering only the current density parallel to the magnetic field, and assuming a quasi-steady state (i.e. neglecting explicit time derivatives), the current density is simply equal to the sum of the Ohmic current density $j_{\text{ohm}} = \sigma E$ and the runaway current density j_{RE} . In a fluid model, the latter can be taken into account by assuming all runaways travel with the speed of light, so that $j_{\text{RE}} = n_{\text{RE}}ec$, where n_{RE} is the runaway electron density.

The part of the equation governing the electron energy density relevant for this work has the form

$$\frac{\partial W}{\partial t} = \nabla \cdot (\mathbf{D} \cdot \nabla W) + P_{\rm ohm} + P_{\rm ion} + P_{\rm rad}.$$
(2.36)

The ion energy density will be neglected for simplicity, which affects the heat capacity at most by a factor of two. The first, diffusive, term on the right hand side represents a cross-field transport of the thermal energy due to perturbations in the magnetic field. The second term on the right hand side is the Ohmic (resistive) heating, $P_{\rm ohm} = jE$. The third term describes the energy lost in order to overcome the electron-ion binding energy during ionization. Finally, the fourth term describes radiative losses, generated by line radiation, recombination radiation due to partially ionized ions and atoms present in the plasma, and by bremsstrahlung. The line radiation comes from excitation-deexcitation processes that give rise to the emission line spectrum of the ions and atoms. Recombination radiation comes from the net energy loss (total loss minus the change in potential energy, which is not lost from the thermal energy content) due to the similar process where a free electron is captured by an ion and relaxed to the ground state.

2.4.3 Electromagnetic fields

The final component needed for a complete plasma model, when using either kinetic or fluid descriptions of the plasma, is a description of the evolution of the electric and magnetic fields. Such a description is given by Maxwell's equations,

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0},\tag{2.37}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{2.38}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t},\tag{2.39}$$

$$\nabla \times \mathbf{B} = \mu_0 \left(\mathbf{j} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right), \qquad (2.40)$$

where ρ denotes charge density. These equations can be used to derive a single equation relating the electric field and the current density. To do this, we make use of the fact that the time variations of interest are slow compared to the typical time scales of electromagnetic waves, so that the displacement current (the second term on the right hand side in the last equation) is negligible. We then take the curl of the third equation and insert the curl of the magnetic field from the fourth equation, yielding

$$\nabla \times (\nabla \times \mathbf{E}) = -\mu_0 \frac{\partial \mathbf{j}}{\partial t}.$$
 (2.41)

The quasi-neutrality condition gives $\rho = 0$, so that $\nabla \cdot \mathbf{E} = 0$. We can then use the identity $\nabla^2 \mathbf{E} = \nabla (\nabla \cdot \mathbf{E}) - \nabla \times \nabla \times \mathbf{E}$ to rewrite the above equation as

$$\nabla^2 \mathbf{E} = \mu_0 \frac{\partial \mathbf{j}}{\partial t}.$$
 (2.42)

This equation indicates the close connection not only between the electric field and the current density, but also between the electric field and the time derivative of the current density. It also contains the earlier mentioned property that the current density cannot change arbitrarily fast (as equation (2.42) puts a constraint on its time derivative). A fast drop in conductivity must therefore be accompanied by an increase in the electric field. These properties are of great importance for the plasma evolution during tokamak disruptions, which is described in detail in the next chapter. 3

Tokamak Disruptions

With the theoretical background from the previous chapter covered, we are now ready to move on to a more detailed description of the main topic of this work, namely disruptions. These events are a type of operational failure in a tokamak, where the plasma suddenly cools. While disruptions present a range of concern already in present-day tokamaks, the larger amount of energy (in the GJ range) released in a reactor scale tokamak disruption make these events the potentially most severe threat to the future of fusion energy based on the tokamak design [26]. A particular concern is the subsequent runaway electron generation, especially due to the so-called runaway avalanche mechanism which exponentiates the number of energetic electrons. This process is exponentially sensitive to the initial current, making it a serious threat to future high-current devices. This chapter starts with a description of the general features of a tokamak disruption, runaway electron mechanisms at play, and potential damage to the device, with an emphasis on the requirements for disruption mitigation in ITER. We give a brief overview of the considered mitigation strategies, and then focus on pellet injection, which is the currently favoured approach and so studied in this work.

3.1 General features

Disruptions are caused by the plasma confinement being subject to a perturbation. These events are closely related to the onset of magnetohydrodynamic (MHD) instabilities [29]. An important quantity in this context is the safety factor, denoted q, defined as the ratio of the number of times a field line revolves in the toroidal direction for every revolution in the poloidal direction. On flux surfaces where q is irrational, the field line never connects to itself, and then a single field line traces out the whole flux surface. This can be used to illustrate the magnetic field geometry with Poincaré-plots, some examples of which are shown in figure 3.1. These plots are obtained by following a magnetic field line and adding a point to the plot every time the field line crosses a given poloidal plane. The resulting illustration of an unperturbed plasma equilibrium is shown in the left panel.

The instabilities are related to perturbations of the plasma on flux surfaces where q is a low-order rational number, referred to as *rational flux surfaces* [29]. On these flux surfaces, the field lines connect with themselves after a few toroidal revolutions, making these surfaces more susceptible to perturbations. The instabilities are often induced by reaching a critical magnidute of a pressure gradient or current density, leading to a growing perturbation of the magnetic field. This can be caused by



Figure 3.1: Poincaré-plots illustrating the stochastisation of a magnetic field in connection to a disruption, (left) unperturbed magnetic field, (middle) creation of magnetic islands at the flux surface where q = 3/2, and (right) stochastisation of the magnetic field due to overlapping of islands centered at different rational flux surfaces. Figure adapted from ref. [30].

e.g an influx of impurities into the plasma. The impurities might originate from an unintended influx of particles from the wall, or from a deliberate injection of impurities attempting to control the energy loss in order to minimize the damage to the device [5, 31]. The latter scenario is the focus of this work, and is introduced in section 3.4.

Perturbations to the magnetic field tend to create island-like magnetic field structures in the vicinity of the rational flux surfaces, centered around the points originally passed by the perturbed field lines. Such structures are illustrated in the middle panel of figure 3.1 for the case with q = 3/2. As such islands are formed and grow on different rational flux surfaces, they might become large enough to overlap with each other. At this point, the magnetic field lines reconnect with each other in a rather chaotic way, causing a stochastisation of the magnetic field [29]. When this happens, the field lines are no longer limited to single flux surfaces, but instead ergodically fill volumes, as illustrated in the right panel of figure 3.1.

As the particles follow the ergodic magnetic field lines across the equilibrium flux surfaces, the particle transport and thermal conductivity across the plasma is greatly increased [26, 31]. If the instability was triggered by an influx of impurities, the impurities also contribute directly to the thermal energy loss by emitting radiation. A major part of the impurity radiation is emitted in the form of line radiation as partially ionized species fall back to lower energy levels after collisional excitation by the surrounding electrons. Another contribution comes from the radiative release of the potential energy change when a free electron recombines with an ion. Moreover, the introduction of high-Z material increases the bremsstrahlung corresponding to the momentum exchange during collisions between charged particles, which might contribute significantly to the thermal energy losses at the high initial temperatures. These mechanisms cause a rapid cooling of the plasma from several keV down to
~ 10 eV over a millisecond time scale. This phase of the disruption is called the *thermal quench*.

As the speed at which particles follow the perturbed field lines is, on average, of the order of the thermal speed, the cross-field transport due to magnetic perturbations decreases with decreasing temperature. In addition, the flux surfaces start to re-heal towards the end of the thermal quench, and the transported losses are therefore greatly reduced after the thermal quench [31]. At this point, the energy balance, and hence the temperature, is essentially determined by impurity radiation and the Ohmic heating by the plasma current.

The cooling leads to a rapid drop in the plasma conductivity, due to the conductivity scaling as $T^{3/2}$, as shown in section 2.3.2. The thermal quench is usually very fast compared to the current diffusion time scale, which is set by the plasma conductivity and the scale and mechanical structure of the device. The current can therefore not change significantly during the thermal quench, so a large electric field must be induced in order to maintain the current. After this, the electric field as well as the current starts to decay by diffusing out of the plasma, a phase called the *current quench* [26]. An example of such an evolution during a simulated ITER disruption is shown in figure 3.2 (data adapted from ref. [19]), where the panels show the evolution of the average temperature and electric field (left), and current (right). The thermal quench and current quench are marked by a red and green shaded area, respectively.

The electric field induced during a disruption is usually well above the critical electric field for runaway electron generation, leading to the conversion of part of the original plasma current into a current carried by runaway electrons [26]. Eventually, the Ohmic current will decay and the runaway current becomes the only current remaining in the plasma. Due to the low collisionality of runaway electrons, this current decays much more slowly than the Ohmic current, forming a *runaway plateau*. This is clearly seen in the blue shaded area of the right panel of figure 3.2. The slow dissipation of the runaway current usually continues until the position control of the plasma is lost and the plasma impacts the wall.

3.2 Vessel loads and mitigation requirements

The various phases of a disruption all pose their own threats to the device, imposing different requirements on a successful mitigation system. We here give a brief overview of the potentially harmful mechanisms at play, and the corresponding requirements for successful disruption mitigation in an ITER-sized tokamak, specified in ref. [5].

The sudden deposition of the released thermal energy content during the thermal quench might cause melting of the plasma facing components if the heat loads are localised. It is therefore necessary for the disruption mitigation system to ensure that a major part of the thermal energy is radiated homogeneously. In ITER, the homogeneously radiated fraction should be larger than 90% of the initial thermal energy content. The time frame between the detection of an emerging disruption and the onset of a naturally occurring MHD-induced thermal quench is expected to be about 20 ms, which sets the required reaction time scale for the disruption mitigation system. Note, however, that any perturbations of the plasma caused



Figure 3.2: Illustration of the general features of a disruption, based on a disruption simulation in an ITER-like plasma (corresponding to Case 1 in ref. [19]). The panels show (left) the average temperature drop and resulting increase and decay of the average electric field, and (right) the current decay. The shaded areas indicate the time spans of the thermal quench (red), current quench (green), and runaway plateau (blue), respectively.

by the mitigation systems can accelerate the growth of MHD instabilities and/or aggravate the runaway electron issue, and the external interventions must therefore be implemented with care.

During the following current quench, the decay of the plasma current and motion of the plasma induces currents in the surrounding structures. The interaction between these currents and the toroidal and poloidal magnetic field gives rise to potentially harmful forces on these structures. As the position control of the plasma is lost and the plasma moves towards the wall, part of the plasma current might flow through the surrounding structures. This part of the plasma current forms a so-called *halo current*, which can also contribute to the vessel loads. The induced currents increase with shorter current quench times, while the risk of substantial halo currents increase with longer current quench times, giving both an upper and lower limit on the acceptable current quench time. In ITER, the current quench time should preferably be longer than 50 ms, with a hard limit at 35 ms, and shorter than 150 ms.

Finally, if a runaway current is formed, a major part of the runaway current can impact the wall when the plasma control is fully lost. This would cause a substantial localised melting, and possibly also damage deeper into the underlying structures. In ITER, the runaway current remaining in the plasma upon wall impact should therefore not exceed 2 MA.

3.3 Runaway electrons

In order for a runaway current to form, there must be a mechanism feeding electrons to the velocity space region above the critical velocity for runaway acceleration, which was introduced in section 2.3.3. Depending on the circumstances, a number of such mechanisms may be present in tokamak disruptions, and these are reviewed in this section.

The runaway generation mechanisms can be divided into two different types: primary generation or seed generation, that is independent of the number of previously present runaway electrons, and secondary generation or avalanche generation, which amplifies an existing runaway seed. An example of the former is the Dreicer mechanism [12]. This mechanism relies upon the fact that the velocity distribution tends to equilibrate collisionally towards a Maxwellian, with a high energy tail above the critical velocity. When this part of the distribution becomes runaways and accelerate to higher energies, the electron bulk will re-equilibrate and "fill out" the depleted tail, resulting in a continuous runaway generation. This can be a major runaway generating mechanism in present day machines [32–34], but is expected to be very small in ITER according to recent simulations, in most cases negligible compared to other mechanisms [16, 19, 35].

The *hot-tail* mechanism of runaway generation occurs because it takes a finite time for the tail of the initially hot Maxwellian velocity distribution to equilibrate to the much lower temperature rapidly obtained by the bulk electrons during the thermal quench [15]. The equilibration of the tail of the distribution is slower than that of the bulk due to the velocity dependence of the collision frequency derived in section 2.3.1. The tail of the distribution might therefore temporarily form a non-Maxwellian electron population at superthermal energies. When the electric field increases at the start of the current quench, a part of this superthermal population may remain at velocities larger than the critical velocity, and in that way become runaways before they have time to thermalise. This mechanism is illustrated in figure 3.3, showing the evolution of the angle-integrated distribution function during a disruption as simulated by the numerical tool described in chapter 4. As opposed to other mechanisms, the hot-tail generation is only present during the initial part of the disruption. The relative importance of this effect is strongly dependent on the details of the temperature drop, primarily the time scale, and can therefore vary over several orders of magnitude depending on the circumstances [36].

In non-nuclear experiments, using pure deuterium plasmas, the Dreicer and hot-tail machanisms are the only primary runaway generation mechanisms. During nuclear operation, however, the power will be generated by fusing deuterium and tritium, with the latter being β^- -radioactive. Part of the energy spectrum of the electron released during the β^- -decay may fall within the runaway region, providing another source of runaway generation [16, 35]. If the post-disruption density is not too large, this mechanism is expected to generate a seed of the order of 1 A in ITER [16, 19]. Despite being seemingly very small, it is calculated to be enough to obtain a final runaway current in the order of several MA due to the avalanche mechanism described below. The runaway generation from tritium decay can however be suppressed by increasing the post-disruption density enough to make the critical energy for runaway generation exceed the energy released during the β^- -decay.

Another runaway generation mechanism present during nuclear operation comes from the activation of the wall due to the bombardment of neutrons released in the fusion reactions. This bombardment makes the wall radioactive, causing it to emit γ -photons. These γ -photons can be Compton scattered against electrons in the plasma, transferring enough energy to an electron to accelerate it over the



Figure 3.3: Illustration of the hot-tail mechanism, showing a representative evolution of the angle averaged electron momentum distribution $\langle f(\bar{p}) \rangle$ as a function of the normalised momentum $\bar{p} = p/(m_e c)$. The vertical dashed lines indicate the runaway threshold momentum (initially outside the scale). The tail of the initially hot Maxwellian distribution (black line) takes a finite time to equilibrate to the dropping temperature of the bulk electrons. As the electric field increases, part of the distribution is therefore "caught" above the runaway threshold momentum, and is accelerated to higher momenta. The data is extracted from the disruption simulation studied later in figure 5.6 a).

runaway threshold [16]. Due to the large energy of the γ -photons, it is practically impossible to increase the density enough to suppress this mechanism, as opposed to the runaway generation due to tritium decay. Moreover, the number of target electrons for Compton scattering increases when the density increases, and hence also the runaway generation rate. Therefore, this mechanism will produce a seed of the order of 0.1-1 A rather independently of the plasma parameters [16, 19].

Finally, the runaway generation by the above mechanisms may be amplified by the *avalanche mechanism*. This mechanism generates runaways through collisions of existing runaways with slower electrons in a way that both electrons have final velocities larger than the critical one after the collision [17, 18, 37]. As the energies of the runaway electrons are much higher than the ionization energy of the ions in the plasma, bound electrons may also contribute to the avalanche process. In fact, it has recently been shown that an increase of the fraction of bound electrons might substantially enhance the avalanche. The reason for this is that the bound electrons contribute to the number of target electrons to practically the same extent as the free electrons, while their contribution to the drag force is smaller than that from free electrons at high electric fields [20].

The avalanche mechanism gives an initially exponential growth of runaway electrons, and is expected to be responsible for the vast majority of the runaway generation in ITER [16, 19]. It has been shown that in circumstances favourable for runaway generation, only a few seed runaway electrons might be enough to generate a

final runaway current of several MA [19]. The exponential growth continues until the runaway current becomes comparable to the total current remaining in the plasma. At this point, the growth of the runaway current significantly reduces the decay rate of the total current, which in turn reduces the induced electric field, as seen near t = 13 ms in figure 3.2, and in that way terminates the runaway generation [20].

3.4 Disruption mitigation strategies

We now move on to the main methods considered in order to achieve the disruption mitigation requirements covered in section 3.2. This is a very active field of research, to which this work aims to contribute. The strategies studied most extensively to date are based on various forms of material injection, and so-called shattered pellet injection has been chosen as the basis for the disruption mitigation system in ITER [8, 21]. This section gives an overview of material injection in the context of disruption mitigation, leading up to the recently suggested two-stage shattered pellet injection scheme that is the focus of this work.

3.4.1 Massive material injection

Massive material injections act to mitigate disruptions in three main ways, corresponding to the requirements on the radiated fraction of the thermal energy, current quench time and runaway avoidance. Suitable materials for radiative dissipation of the thermal energy are noble gases such as neon or argon [5]. Also note that as long as some amount of the impurity is present, the radiation can be further enhanced by increasing the electron density by injecting e.g. hydrogen species, since the collisional excitation rate is proportional to the electron density.

The quantity and composition of injected material can also be used to regulate the post-disruption temperature. As mentioned in section 3.1, this temperature is roughly given by an equilibrium between the impurity radiation and the Ohmic heating. Since the current quench time is proportional to the conductivity, which in turn scales as $T^{3/2}$, the temperature essentially determines the current quench time. A suitable temperature to give an acceptable current quench time in ITER lies in the 5-10 eV range [20].

Finally, the injected material might reduce the runaway generation due to the resulting electron density increase, leading to an increase in the critical electric field according to equation (2.17). The usefulness of this method to reduce the runaway generation has however recently been questioned [19], for reasons detailed below in section 3.4.2.1.

With the above general background about the purpose of massive material injection in mind, we now turn to the different methods proposed to deliver the injected material in the following three subsections.

3.4.1.1 Gas injection

The most straightforward injection method is *massive gas injection*, where the injected material is simply released in gaseous form from a pressurised vault [5]. As

the simplest proposed method to implement, it is the one most studied experimentally, and therefore also theoretically. The greatest advantage with this method is its simplicity, while the assimilation efficiency of the released gas is found to be rather poor in practice on larger machines, despite promising results on smaller machines [38, 39]. Much of the material has been observed in both experiments and simulations to be stopped at the plasma edge and only slowly mix with the rest of the plasma. This can be understood as a result of the fact that as soon as the gas particles ionise they become confined by the magnetic field, which together with the pressure of the pre-existing plasma restricts the gas penetration. The resulting strong cooling at the edge may induce substantial MHD activity, which may help to speed up the inward transport of injected ions but also lead to an unwanted increase in the conducted heat losses.

3.4.1.2 Pellet injection

Another method that gives a faster, more efficient delivery of material to the plasma core, is to inject the material in the form of solid cryogenic pellets [5]. In this way, the material travels through the plasma in a neutral, solid form, while being continuously ablated by the hot background plasma, depositing material along the trajectory. The pellets are typically accelerated by a propellant gas and reach speeds of around 300-600 m/s. This is similar to many typical gas sound speeds, and therefore the arrival time of the pellets at the plasma is not significantly different compared to a gas injection from the same location. Note that injection of pellets (though typically slightly smaller and of hydrogen isotopes) is regularly performed on existing machines also for other purposes than disruption mitigation, such as to fuel the plasma, regulate instabilities and for diagnostic purposes [40, 41].

3.4.1.3 Shattered pellet injection

There are, however, a number of disadvantages associated with disruption mitigation by pellet injection. Depending on the speed at which the pellet travels and the state of the plasma, the pellet may pass through the plasma without depositing all of its material. Besides making the injection less effective, the remaining pellet might damage the wall upon impact [5]. This problem is particularly relevant when a pellet is injected after the plasma has already been cooled in the course of the thermal quench. One way to address this issue is to use *Shattered Pellet Injection (SPI)* [42]. In an SPI, the pellet is shattered against a tilted plate before entering the plasma. The number of shards into which the pellet is shattered may be controlled (to some extent) by varying the speed and impact angle on the shattering plate. Increasing the number of shards (for a fix total amount of pellet material) increases the ablation rate, hence reducing the amount of material passing through the plasma without ablating. The increased ablation may be understood by the fact that shattering the pellet increases the total contact area with the plasma. Any leftover material striking the wall will also be spread over a larger area, reducing the risk of damaging the wall. The initial spread of the deposited material is also increased, lowering the risk of local peaking in the radiative heat loads before the material has homogenised through the plasma.

Because of these advantages, shattered pellet injection has been chosen as the basis of the ITER disruption mitigation system [8, 21]. The current design of the ITER disruption mitigation system, described in ref. [21], includes 24 injectors distributed over three toroidal locations. Each injector can inject cylindrical pellets with a diameter of 28.5 mm and a height of 57 mm, consisting of deuterium and/or neon, or a mixture thereof. Such pellets contain of the order of 10^{24} atoms, with some dependence on their composition. The pellets will be accelerated either by a pressurized gas or an electromagnetic driven punch system. For the lighter deuterium pellets, speeds of up to 800 m/s are expected to be achievable, while the maximum speed of the heavier neon pellets is expected to be about 200 m/s. However, the design and operation parameters of the ITER disruption mitigation system, such as the pellet composition, the number of pellets and their particle contents, the degree of shattering, timing aspects etc. remain open questions. These questions are currently being addressed by major simulation efforts [16, 19, 24], including this work. The current state of research regarding these questions is the subject of section 3.4.2.

3.4.2 Injection schemes

A particular difficulty for the disruption mitigation system is that the various requirements of the disruption mitigation are to some extent contradictory [5]. For the mitigation of the thermal loads, it would be beneficial to have an early large injection of strongly radiating material such as argon or neon. Such injections could also substantially increase the electron density, which would contribute to the reduction of the runaway generation. Large amounts of argon or neon might however result in a post-thermal quench temperature too low to give an acceptable current quench time. Moreover, they might also increase the runaway seed generation from the hot-tail mechanism as well as enhance the subsequent runaway avalanche due to the presence of bound electrons. This section gives an overview of the current state of research regarding suggested compromises attempting to circumvent these issues.

3.4.2.1 Mixed deuterium-impurity injection

The currently envisaged compromise is to inject large amounts of deuterium combined with a trace amount of argon or neon [21, 42]. The role of the deuterium would be to provide a source of electrons in order to limit the runaway generation and to enhance the radiation efficiency of the neon. The role of the neon would be to radiatively dissipate the thermal energy and to set a post-thermal quench equilibrium temperature within an acceptable range.

The density required to give a sufficient runaway reduction is however still under investigation. The electron density required to make the critical electric field larger than the maximum electric field expected to be induced during disruptions in the nuclear operation phase of ITER is of the order of 10^{22} m^{-3} [5]. This density is about two orders of magnitude larger than the typical density during normal operation. On the other hand, it was found in ref. [16] that a density of $4 \cdot 10^{21} \text{ m}^{-3}$ is sufficient to reduce the runaway generation to an acceptable level. Moreover, it was found in ref. [19] that the reduction might only be present for a certain range of injected impurity and deuterium densities. When the injected impurity and/or deuterium densities increase, the post-disruption temperature decreases, and if the post-disruption temperature becomes close to 1 eV, all species present in the plasma, including the hydrogen isotopes, start to recombine. The resulting increase in the fraction of bound electrons greatly enhances the runaway avalanche, as mentioned in section 3.3. In addition, the runaway generation further increases due to the increase in the electric field resulting from the lower temperature. The trend of reducing runaway generation was therefore found in ref. [19] to turn at deuterium densities of the order of 10^{21} m⁻³ for an ITER plasma in the nuclear operation phase. The calculated minimum runaway current (within the range of parameters that also gave an acceptable current quench time) was 3.7 MA. It was however noted that these values are sensitive to details in the model of the radiation losses that contain significant uncertainties, such as the opacity of the plasma to the emitted line radiation, requiring further investigation.

3.4.2.2 Multiple-stage injection

A recently suggested improvement to the above scheme is to divide the injection into two stages following rapidly after each other [24]. The first stage would then deliver the deuterium, and the second deliver the neon or argon. The aim of such a scheme would be to first cool the plasma by dilution down to the 100-1000 eVrange by the pure deuterium injection, without perturbing the plasma pressure or current density enough to significantly accelerate the growth of MHD instabilities. The plasma would then be left at this temperature for a few milliseconds to let the full distribution equilibrate to a Maxwellian at this temperature. A final radiative thermal quench would then be triggered by injecting the argon or neon content. The intermediate equilibration of the distribution could potentially produce a significant reduction of the hot-tail runaway generation. The radiated fraction of the thermal energy could also be significantly increased, as the magnetic perturbations would not become significant until the comparatively low temperature makes conducted losses subdominant to radiation losses, as discussed in section 3.1. It was indicated in ref. [24] that it is possible to cool an ITER-like plasma by dilution down to ~ 100 eV without immediately triggering an MHD-induced thermal quench. The runaway dynamics and radiation characteristics of such a two-stage SPI scheme have however not been thoroughly studied before, so this scheme is the main focus of chapter 5 of this work.

3.5 Pellet injection physics

In order to assess the performance of the injection schemes described in section 3.4.2, an understanding of the underlying physics of pellet injection is essential. Before going into the details of the disruption model used in this work, presented in chapter 4, we therefore conclude this chapter with an overview of the background theory of pellet injection physics. The process is divided into two steps: the ablation of the pellet (or pellet shards) by the hot background plasma, covered in section 3.5.1,

and the subsequent homogenisation and equilibration over the flux surfaces of the ablated material, covered in section 3.5.2.

3.5.1 Pellet ablation

In essence, the number of particles ablated from the pellet during a given time interval is determined by the ratio of the heat flux reaching the pellet surface during this time interval and the sublimation energy per particle. However, on a very fast time scale after the pellet is exposed to the plasma (of only a few μ s) the finite flow of ablated material away from the pellet produces a dense cloud around the pellet that shields it from the heat flux from the plasma [43]. The pellet ablation can therefore be regarded as a self-regulatory process, balancing the heat flux from the plasma and the resulting build-up of the shielding cloud. The resulting quasi-stationary cloud size and ablation rate are such that only the ablation energy necessary for maintaining the cloud size reaches the pellet surface.

Directly above the pellet surface, the cloud is neutral and close to spherically symmetric. This neutral cloud typically has a thickness of the order of 1 cm, and a particle density of $10^{25} - 10^{26}$ m⁻³. The pellet (or pellet shard) itself typically has a thickness of the order of a few millimeters. At the edge of the neutral cloud, the material density falls and the temperature increases enough to begin to ionize the ablated material. The material takes the form of a confined cool plasma, whose subsequent expansion is therefore mostly aligned with the field lines.

There are three main mechanisms involved in the shielding of the pellet due to the neutral cloud, although usually with quite different importance [43]. With decreasing importance, these mechanisms are referred to as gas dynamic shielding, electrostatic shielding and diamagnetic shielding. The gas dynamic shielding refers to the deposition of incident plasma energy by the collisional interaction with the neutral gas in the closest proximity of the pellet. Collisions between the hot incident electrons and the new cold ablated plasma beyond the neutral cloud also contribute to some extent, but gives a typically much smaller contribution [44, 45]. The energy deposition in the cloud occurs due to scattering, heat transfer to the cloud particles, and ionization and excitation of ions and atoms in the cloud, which then dissipate the energy by radiation [43, 46]. It has however been estimated that, due to its relatively high density, the cloud might be substantially opaque to the resonant lines. A significant fraction of the radiation might therefore be trapped within the cloud, and the corresponding energy also contribute to the cloud heating [47, 48].

Electrostatic shielding results from the different mobility of ions and electrons [43]. Electrons in the background plasma initially flow into the ablation cloud much faster than the ions, due to their lower mass. This gives rise to a difference in charge between the ablation cloud and the surrounding plasma, and the corresponding electrostatic potential difference limits the heat flux into the ablation cloud.

Diamagnetic shielding occurs as there is a finite timescale for the magnetic field diffusion into the cloud [43]. The gyration of the ionized particles in the ablation cloud opposes the background magnetic field, hence the name diamagnetic shielding. The heat flux from the background plasma is guided by the magnetic field lines, so deflection of the magnetic field leads to a deflection of the heat flux around the ablation cloud. The importance of this effect can be characterised by the ratio of the flow speed of material away from the pellet and the speed of the diffusion of the magnetic field into the neutral cloud, which is not expected to be very large due to the comparatively low conductivity of the neutral cloud [44].

In reality, some further complications compared to the picture described above arise. One such complication is that as the pellet moves through the plasma, it can periodically cross the boundary of the ablation cloud surrounding it and establishes a new one upon the direct exposure to the background plasma [41]. This periodicity might explain the striations in the deposition profile sometimes observed in experiments. Another complication is that the ablation is not completely symmetric around the pellet parallel to the field lines [41, 43]. Reasons for this includes the difference in the heat flux parallel and anti-parallel to the electric field. This makes the ablation cloud thicker in the direction anti-parallel to the electric field. Another source of cloud asymmetry is irregularities in the pellet shape. Asymmetric ablation might give rise to a "rocket-like" effect, propelling the pellet in the opposite direction to the increased ablation. In many theoretical models, however, these subtleties are not accounted for.

3.5.2 Pellet material deposition

Once the material is ablated and ionized, it begins to homogenise over the flux surfaces and the pressure and temperature start to equilibrate with the background plasma. In present day machines the homogenisation process takes place over a time scale of the order of 0.1-1 ms [40]. Apart from the increased density, the equilibrated post-injection plasma has also been cooled by dilution of the initial thermal energy and the energy consumed by sublimation and ionization.

There are two main processes responsible for the homogenisation. The first is the excess pressure of the cold ablated plasma driving its expansion along the magnetic field. As a field line covers a whole flux surface (except at rational flux surfaces), the expansion along the field lines eventually leads to the material being homogenised over the entire flux surface. This mechanism alone has however been shown to give about an order of magnitude slower equilibration than the 0.1-1 ms time scale observed in experiments [40].

An additional mechanism to consider in describing the homogenisation is caused by the potential difference between the channel of ablated material and the background plasma, again arising due to the much larger mobility of electrons compared to ions. As the channel of ablated material is heated while still having a much larger density than the background plasma, the net flow of negative charge will go from the channel of ablated material into the background plasma. The gradient along the tokamak minor radius of this potential gives rise to an electric field along the minor radius, which in turn causes an $E \times B$ -drift in the poloidal direction. Conservation of momentum on the flux surface where the channel of ablated material resides then gives rise to a poloidal rotation of the plasma in the direction opposite to the poloidal $E \times B$ -drift of the ablated material. The gradient of this rotation along the minor radius, together with the variation of the background magnetic field direction, gives rise to a poloidal stretching of the ablated material. As the ablated material is thinned out, the collisional interaction with the background plasma becomes more effective, and eventually the ablated material equilibrates with the background plasma. The above mechanisms together have been shown to give rather accurate reproductions of the time evolution of the background plasma following fueling pellet injections in experiments [40].

Finally, yet another $E \times B$ -related drift might be present that transports the ablated material across the flux surfaces [49]. The limited length along the toroidal field line of the initial expanding channel of ionised, ablated material, re-introduces the issue of vertical charge separation in the channel by the grad B drift. Remember this issue was described in section 2.2.2, occurring in the whole tokamak in the absence of a poloidal magnetic field. This gives rise to a vertical electric field in the vicinity of the ablated material, causing an $E \times B$ drift of the ablated material along the major radius towards the low field side. This effect results in a shift of the final deposition profile compared to the deposition profile immediately following the ablation.

3. Tokamak Disruptions

4

Disruption model

Having acquired a basic knowledge of the characteristic features of disruptions, and the proposed SPI schemes to mitigate their impact, we are now in a position to describe the model used in this work to simulate such scenarios. Following the SPI, the pellet shards begin to ablate when they are exposed to the hot background plasma. Once the pellet material is deposited, the plasma density increases and the temperature drops. The cooling occurs partly due to the dilution of the thermal energy, and partly due to energy losses through radiation as well as transport stemming from the perturbation of the magnetic field, as detailed in section 3.1.

A number of components are required to build a model for the pellet material deposition through the final current decay, and they are described in the sections of this chapter. The model for the evolution of the pellet and the resulting density increase is described in section 4.1, and the model describing the subsequent plasma cooling is described in section 4.2. The cooling of the plasma is accompanied by a drop in the conductivity, leading to the induction of a strong electric field, after which the plasma current starts to decay. The strong electric field might however lead to the generation of a substantial runaway current, as was also described in section 3.1, resulting in an incomplete current decay. The evolution of the electric field is governed by the equations described in section 4.3, and the models governing the runaway generation resulting from the induced electric field are described in section 4.4. Finally, the numerical tool DREAM [25], which solves the set of equations listed in this chapter starting from a set of given SPI parameters and pre-disruption plasma parameters, is described in section 4.5.

4.1 Shattered pellet injection

An SPI starts with a pellet being accelerated and then shattered against a tilted plate, resulting in a plume of pellet shards of different sizes and velocities entering the plasma. As the shards travel through the plasma, they provide a set of localised, moving sources of particles. In the 1D model used in this work, assuming an instantaneous homogenisation over the flux surfaces, such a particle source translates to a source acting on the densities at the local flux surface. The strength of the density sources are determined by the ablation rate of the shards, the size of the flux surfaces, and a kernel function defining the radial spread of the deposited material around the shards. The ablation rate must be calculated self-consistently with the buildup of the ablation cloud, shielding the pellet from further ablation, as described in section 3.5.1. This requires a model for the shielding of the heat flux from the ambient plasma in the ablation cloud.

Here, we employ the gas dynamic shielding description, introduced in section 3.5.1. Furthermore, the neutral ablation cloud directly above the pellet surface. where most of the shielding takes place, is assumed to be spherically symmetric. Thus, the problem is reduced to one dimension, with the radius from the center of the pellet as the only coordinate. The energy is primarily deposited in the cloud by the light electrons of the bulk plasma, which stream rapidly along the field lines to intersect the pellet. The gas dynamic shielding model determines the reduction of the incident electron heat flux on the way through the shielding cloud, in terms of the density and energy distribution of the bulk plasma electrons. This requires expressions for the electron energy reduction due to absorption and scattering in the neutral cloud, for a given energy of the incident electron. The model moreover imposes a boundary condition at the pellet surface, such that a quasi-stationary ablation process is achieved, where the energy reaching the pellet surface is just sufficient to provide the sublimation energy required to maintain the cloud. This boundary condition is determined by the sublimation energy and flow properties at the solid-gas interface.

Analytical solutions for the pellet ablation rate within the above setup have been determined in the literature. The original treatment was given in ref. [44] for pure hydrogen pellets. The energy distribution of the incident electrons from the bulk plasma is typically approximately Maxwellian, which in ref. [44] was approximated by a single mono-energetic beam with effective energy per particle $\mathcal{E}_{in} = 2T_M$, equal to the ratio of the unidirectional heat and particle flux. Here, T_M denotes the temperature of the Maxwellian population. In our notation, the subscript M is included to distinguish the Maxwellian part of the electron population from the superthermal electron population that might form during a disruption, as described in section 3.3. The model for the electron energy absorption and scattering consisted of empirical expressions specific for hydrogen isotopes. Due to the small sublimation energy for hydrogen, the boundary condition at the pellet surface was set by assuming the heat flux and particle energy to be negligible there. The ablation rate thus derived is called the *Neutral Gas Shielding (NGS)* model.

This model was later improved upon, by accounting for the full Maxwellian energy distribution of the incident particles instead of a mono-energetic beam in ref. [48]. A treatment of pellets consisting of higher atomic numbers was made in ref. [50], accounting for the Z-dependence of the electron energy absorption and scattering, and the effect on the boundary conditions resulting from the higher sublimation energy. The most up to date version of the NGS model, accounting for the full Maxwellian electron momentum distribution with a heat flux reduction model valid for pellets consisting of a mix of neon and deuterium was presented in ref. [51]. This analytical model was recently benchmarked against a numerical time-dependent 3D Lagrangian flow model in ref. [52], showing agreement within a few percent when using similar geometrical approximations. Despite its simplicity, various versions of the NGS model has been shown to produce simulated density profiles that agree with experiments within a few tens of percent, for both fueling pellets [43] and massive SPI injections [53]. Therefore, we adopt the NGS model of ref. [51] here, for these first SPI disruption simulations with DREAM. The technical aspects needed to describe the SPI scenario are presented below in more detail. These include the size and velocity distributions of the shattered pellet shards (with parameters to be optimised), and the explicit form of the ablation rate. Finally, the choices regarding the translation from a particle source to a density source are discussed further.

4.1.1 Shattering

We assume the pellets to be shattered into $N_{\rm s}$ shards, which are approximated as spherical, with sizes described by an equivalent radius $r_{p,k}$, with $k = 1, ..., N_{\rm s}$. The shard radii are assumed to be drawn randomly from the distribution with probability density

$$P(r_{p,k}) = k_p^2 r_{p,k} K_0(k_p r_{p,k}), \quad k_p = \left(\frac{N_{\text{inj}}}{6\pi^2 n_p N_s}\right)^{-1/3}, \quad (4.1)$$

where K_0 is the zeroth modified Bessel function of the second kind, n_p is the number density of the solid pellet material, and N_{inj} is the total number of injected atoms. This form of the shard size distribution reflects the fact that the pellet is initially mostly broken by shear stresses into a saucer-like structure [54]. If the further breakup of the thin layers thus formed is approximated as a division into a large number of rectangles, separated by randomly and independently distributed perpendicular straight lines, the form of the shard size distribution given by equation (4.1) is obtained [55]. This distribution of shard sizes has recently been used in several other SPI studies [56–58].



Figure 4.1: Illustration of an SPI injection, as modeled here, defining the geometry and parameters used to describe the motion of the pellet shards.

Once shattered, the shards are assumed to travel with constant velocities $v_{p,k}$ in a 2D-plane, starting at the shattering point (x_0, y_0) , as illustrated in figure 4.1;

$$\boldsymbol{x}_{p,k}(t) = (x_{p,k}(t), y_{p,k}(t)) = (x_0 - v_{p,k} \cos \alpha_{p,k} t, y_0 + v_{p,k} \sin \alpha_{p,k} t), \quad (4.2)$$

with the origin of the (x, y)-coordinate system positioned at the plasma center. The speeds $v_{p,k}$ and angles $\alpha_{p,k}$ to the horizontal plane are chosen from uniform random distributions within $\langle v_p \rangle \pm \Delta v_p$ and $\pm \Delta \alpha_p$, respectively. The parameters $x_0, y_0, \langle v_p \rangle$, Δv_p and $\pm \Delta \alpha_p$, as well as N_{inj} and N_s , are considered controllable free parameters. In reality, the control of these parameters is achieved by adjusting the impact speed and angle of the pellet on the shattering plate (see section 3.4.1.3). This is likely to result in correlations between the pellet size and the distribution of shard sizes and velocities. However, in lieu of a self-consistent model to capture these correlations, the injection parameters are assumed to be independently variable in this work.

4.1.2 Ablation

We characterise the ablation rate by the time derivatives $\dot{r}_{\mathrm{p},k}$ of the shard radii. The expression used for $\dot{r}_{\mathrm{p},k}$ is based on the updated NGS model presented in ref. [51]. Expressed in terms of the unidirectional incident heat flux q_{in} carried by the bulk plasma electrons and their effective energy $\mathcal{E}_{\mathrm{in}}$, this model gives

$$\dot{r}_{p,k} = -\lambda(X) \left(\frac{q_{\rm in}}{q_0}\right)^{1/3} \left(\frac{\mathcal{E}_{\rm in}}{\mathcal{E}_0}\right)^{7/6} \left(\frac{r_{p,k}}{r_{p0}}\right)^{4/3} \frac{1}{4\pi r_{p,k}^2 \rho_{\rm dens}}.$$
(4.3)

Here, the normalising radius, heat flux and effective energy are $r_{p0} = 2 \text{ mm}$, $q_0 = n_0 \sqrt{2T_0^3/(\pi m_e)}$ and $\mathcal{E}_0 = 2T_0$, with the representative temperature and density $T_0 = 2000 \text{ eV}$ and $n_0 = 10^{20} \text{ m}^{-3}$, respectively. The solid mass density of the pellet is denoted ρ_{dens} . The dependence on the deuterium-neon composition is accounted for by the factor

$$\lambda(X) = [27.0837 + \tan(1.48709X)]/1000 \text{ kg/s},$$

where $X = N_{\text{D}_2}/(N_{\text{D}_2} + N_{\text{Ne}})$ is the deuterium fraction, N_{D_2} is the number of deuterium molecules and N_{Ne} is the number of neon atoms in the pellet. Note that X is defined in terms of the number of deuterium molecules in the pellet, while $N_{\text{inj}} = 2N_{\text{D}_2} + N_{\text{Ne}}$ is the number of injected atoms.

The heat flux and effective energy are calculated from a general electron momentum distribution function, f, according to

$$q_{\rm in} = \frac{1}{4} \int m_{\rm e} c^2 (\gamma - 1) v f \,\mathrm{d}\boldsymbol{p} \tag{4.4}$$

and

$$\mathcal{E}_{\rm in} = \frac{2}{n_{\rm free}} \int m_{\rm e} c^2 (\gamma - 1) f \,\mathrm{d}\boldsymbol{p}. \tag{4.5}$$

We note that equation (4.3) was derived assuming a Maxwellian electron momentum distribution, and is not strictly applicable for a general distribution. However, it may be assumed to be sufficiently accurate for the small deviations from a Maxwellian present in the early stages of the disruption while the shards are still ablating, that is, before a substantial runaway acceleration has occurred. The total free electron density is $n_{\rm free} = \int f d\mathbf{p}$, c is the speed of light, and γ is the Lorentz factor. The factor 1/4 in equation (4.4) converts the isotropic heat flux to the average unidirectional heat flux facing the pellet shards, and is strictly valid for a Maxwellian distribution [59], whilst $\mathcal{E}_{\rm in}$ reduces to $2T_{\rm M}$ for completely Maxwellian electrons.

4.1.3 Material deposition

When the shards are ablated, the ablated material ionizes, and is therefore largely confined by the magnetic field to the flux surfaces in the vicinity of the position of the pellet shard where the ablation took place. As mentioned in section 3.5.2, the homogenization of the ablated material over the flux surfaces, as well as the temperature and pressure equilibration with the background plasma, occurs over a time scale ≤ 1 ms. This time scale is comparable to the thermal quench time scale, as well as the time between the arrival of the first and last shard on a given flux surface, while being relatively fast compared to the current quench time scale. The homogenization and equilibration of the ablated material is therefore approximated here to take place instantaneously, an assumption also made in other recent SPI studies [24, 53, 57, 58]. The impact of this assumption is discussed in Appendix C. Similarly, all other quantities involved, except for the pellet shard positions and radii, are assumed to be constants over the flux surfaces. The flux surfaces are assumed to have concentric circular cross-sections and are labeled with their minor radius r. Moreover, we employ the large aspect ratio limit $R_0/a \gg 1$, where R_0 and a are the major and minor radii of the plasma, respectively, so that the spatial geometry becomes cylindrically symmetric.

The homogenized ion density increase on the flux surface with radius r due to the ablation of the pellet material is given by

$$\left(\frac{\partial n_{ij}}{\partial t}\right)_{\rm SPI} = -f_{ij} \sum_{k=1}^{N_{\rm s}} \frac{4\pi r_{p,k}^2 \dot{r}_{p,k} \rho_{\rm dens} N_{\rm A}}{\mathcal{M}} H(r,\rho_{p,k}),\tag{4.6}$$

where n_{ij} is the density of charge state *i* of ion species *j*, and the factor f_{ij} denotes the particle fraction of the ablated material that is deposited to n_{ij} . The pellet molar mass is denoted by \mathcal{M} , and $N_{\rm A}$ is the Avogadro number. The radial distribution of the homogenized density increase is described by the factor $H(r, \rho_{p,k}) = h(r, \rho_{p,k})/A_{\rm fls}(r)$, where $h(r, \rho_{p,k}) dr$ describes the fraction of the material deposited at a radius between *r* and *r* + d*r* ablated from a pellet at radius $\rho_{p,k} = \sqrt{x_{p,k}^2 + y_{p,k}^2}$. Here, $A_{\rm fls} = 4\pi^2 r R_0$ is the area of the flux surface at radius *r*.

The width of the volume within which the ablated material is deposited may be approximated by the width of the shielding cloud around the pellet, with a radius denoted $r_{\rm cld}$. To account for the finite width of this cloud, previous studies have used a Gaussian deposition kernel $h \propto \exp\left[(r - \rho_{p,k})^2/r_{\rm cld}^2\right]$ [24, 56]. The calculations made in e.g. ref. [46] show that a realistic value for $r_{\rm cld}$ is of the order of 1 cm. However, in many of the cases studied in this work, we are required by computational feasibility to have a radial resolution that is significantly larger than such a width of the pellet cloud, and the radial resolution then becomes the limiting length scale for the spread of the ablated material. Moreover, since the numerical tool used in this work, described in section 4.5, uses an implicit time-stepping carried out by a Jacobian-based Newton solver, the computational expense increases rapidly with the number of non-zero Jacobian-elements introduced by spreading the ablated material over several flux surfaces. For these reasons, we use a delta function deposition kernel, $h = \delta(r - \rho_{p,k})$ which, when discretized in time, translates to a uniform distribution over the distance traveled during the current time step, as detailed in Appendix A. The impact of the radial grid step size and the use of a Gaussian deposition kernel with a finite width is investigated in Appendix C. These calculations verify that the width of the deposition kernel has only a formal impact on the final density profile, even when the radial grid step size is smaller than the pellet cloud width.

At the relatively high plasma temperatures into which the pellets are injected, the fast ionization process of the lower charge states introduces a problematic bottleneck in the need for time resolution. To circumvent this issue, we deposit the ablated material directly to the equilibrium distribution of charge states associated with the local density and temperature. With ϕ_j denoting the particle fraction of the pellet material consisting of species j, the quantity f_{ij} appearing in equation (4.6) becomes $f_{ij} = \phi_j n_{ij}^{\text{eq}}/n_{\text{tot},j}$, where the equilibrium distribution of charge states is calculated according to

$$R_{ij}n_{i+1,j}^{\rm eq} - I_{ij}n_{ij}^{\rm eq} = 0, \quad i = 0, 1, ..., Z - 1,$$

$$\sum_{i} n_{ij}^{\rm eq} = n_{\rm tot,j}, \quad i = 0, 1, ...Z.$$
(4.7)

The total density of species j is denoted by $n_{\text{tot},j}$, and $I_{ij}(T_M, n_M)$ and $R_{ij}(T_M, n_M)$ are the ionization and recombination rates, respectively, obtained from the OpenADAS database [60]. The density n_M of the Maxwellian population is taken to be the density of electrons with momenta lower than a threshold p_{hot}^{-1} , described in section 4.2, i.e. $n_M = \int_{0 . The local increase in the electron density due to the$ deposition of the pellet material is determined by the quasi-neutrality condition, $<math>n_{\text{free}} = \sum_{ij} Z_{ij} n_{ij}$, where Z_{ij} is the charge number of charge state i of ion species j. The new electrons are added to the momentum distribution function at zero momentum, as described in section 4.4.1, and are thus added to the density n_M of the Maxwellian population.

4.2 Particle and energy balance

When the cold pellet material is deposited, the hot plasma is initially cooled simply by dilution of the thermal energy on the local flux surface over the new particles. Thereafter, the partially ionized particles in the deposited material dissipate the thermal energy as radiation, either directly through line radiation or indirectly by ionization, resulting in a final radiative loss during recombination. Moreover, the introduction of high-Z impurities is expected to trigger or accelerate the growth of MHD instabilities resulting in a stochastisation of the magnetic field. This stochastisation increases the transport of thermal energy out of the plasma, as discussed in section 3.1. The thermal energy is further affected by local Ohmic heating from the thermal part of the plasma current and, at high temperatures, bremsstrahlung losses.

When evaluating the energy transport in the plasma, as well as the current density described in sections 4.3 and 4.4, we treat the Maxwellian electrons and

¹Although this part of the momentum distribution is not strictly Maxwellian, the collision operator is linearised around a Maxwellian with density $n_{\rm M}$, as described in section 4.4.1, hence the subscript M.

the superthermal electrons separately. The separation is made by introducing a momentum p_{hot} , and considering electrons with a momentum larger than p_{hot} as superthermal. To increase the computational efficiency, we further divide the superthermal electrons into two parts referred to as the hot population and the runaway population, respectively, separated by a momentum p_{RE} . The hot population is characterized by the part f_{hot} of the distribution function with $p_{hot} ,$ $while the runaway electrons are characterized by a density <math>n_{RE}$, and are approximated to travel with the speed of light. Typically, we set $p_{hot} \approx 10p_{th}$, where p_{th} is the thermal momentum (making p_{hot} a time-varying boundary), and p_{RE} equal to a few times m_ec . This choice of p_{RE} corresponds to a speed within a few percent of c, justifying the approximation of the speed of the runaway population. We note that some of the electrons in the hot population might formally satisfy the definition of a runaway electron if the critical runaway momentum lies between p_{hot} and p_{RE} , but these electrons will accelerate to a momentum higher than p_{RE} in a comparatively short time scale.

We now turn to the model used for the evolution of the energy density of the Maxwellian electrons, $W_{\rm M} = 3n_{\rm M}T_{\rm M}/2$, and the ion charge state densities n_{ij} . The evolution of the non-Maxwellian electrons, including the evolution of $f_{\rm hot}$ and $n_{\rm RE}$, will be treated in sections 4.4.1 and 4.4.2. The evolution of $W_{\rm M}$ on the local flux surface is governed by

$$\frac{\partial W_{\rm M}}{\partial t} = \sigma_{\rm M} E_{||}^2 - n_{\rm M} \sum_{ij} n_{ij} [L_{ij}(T_{\rm M}, n_{\rm M}) + E_{ij}^{\rm ioniz} I_{ij}(T_{\rm M}, n_{\rm M})] \\
+ \left(\frac{\partial W_{\rm M}}{\partial t}\right)_{\rm ioniz}^{\rm abl} + \frac{1}{r} \frac{\partial}{\partial r} \left[r D_W \frac{\partial T_{\rm M}}{\partial r} \right] + \int_{p_{\rm hot}$$

The first term on the right hand side corresponds to the Ohmic heating. The conductivity $\sigma_{\rm M}$ is calculated using the expression derived in ref. [61], with relativistic effects taken into account;

$$\sigma_{\rm M} = \bar{\sigma} \frac{4\pi\epsilon_0^2 T_{\rm M}^{3/2}}{Z_{\rm eff}\sqrt{m_e e} \ln \Lambda_0},\tag{4.9}$$

where $\bar{\sigma}(T_{\rm M}, Z_{\rm eff})$ is calculated by interpolation of the values tabulated in ref. [61]. This expression accounts for the contribution from both the test particle and field particle part of the Coulomb collision operator, and is discussed further in section 4.3. Here we have also introduced the effective charge $Z_{\rm eff} = \sum_{ij} Z_{ij}^2 n_{ij}/n_{\rm free}$, the dielectric constant ϵ_0 , the elementary charge e, the thermal Coulomb logarithm $\ln \Lambda_0 = 14.9 - 0.5 \ln (n_{\rm M}/10^{20}) + \ln (T_{\rm M}/10^3)$ [29], and the electric field parallel to the magnetic field lines $E_{||}$. The component of the electric field perpendicular to the magnetic field lines is neglected in agreement with the assumption of cylindrical symmetry.

The second term corresponds to the line radiation and ionization losses. The line radiation rates $L_{ij}(T_{\rm M}, n_{\rm M})$ are taken from the OpenADAS database [60], and the ionization energies $E_{ij}^{\rm ioniz}$ are taken from the NIST database². Note that ionization

²https://physics.nist.gov/PhysRefData/ASD/ionEnergy.html

losses are still present when the net electron density does not increase, as in that case the ionization and recombination processes are still active although they balance each other. The change in the local charge state distribution of all ions due to ionization and recombination is calculated by the time dependent rate equations

$$\left(\frac{\partial n_{ij}}{\partial t}\right)_{\text{ioniz}} = I_{i-1,j}n_{i-1,j}n_{\mathrm{M}} - I_{ij}n_{ij}n_{\mathrm{M}} + R_{i+1,j}n_{i+1,j}n_{\mathrm{M}} - R_{ij}n_{ij}n_{\mathrm{M}}.$$
 (4.10)

Thus, the total evolution of the charge state densities is given by

$$\frac{\partial n_{ij}}{\partial t} = \left(\frac{\partial n_{ij}}{\partial t}\right)_{\text{ioniz}} + \left(\frac{\partial n_{ij}}{\partial t}\right)_{\text{SPI}}.$$
(4.11)

The corresponding evolution of the electron density is determined by the quasineutrality condition.

The third term in equation (4.8) accounts for the energy required to ionize the ablated pellet material to the equilibrium charge state distribution described in section 4.1. This gives an additional ionization loss according to

$$\left(\frac{\partial W_{\rm M}}{\partial t}\right)_{\rm ioniz}^{\rm abl} = -\sum_{ij} \Delta E_{ij}^{\rm binding} f_{ij} \sum_{k=1}^{N_{\rm s}} \frac{4\pi r_{p,k}^2 \dot{r}_{p,k} \rho_{\rm dens} N_{\rm A}}{\mathcal{M}} H(r, \rho_{p,k}), \qquad (4.12)$$

where $\Delta E_{ij}^{\text{binding}} = \sum_{0}^{i-1} E_{ij}^{\text{ioniz}}$ is the total energy required to ionize an atom of species j from neutral to charge state i. The assumption of instantaneous homogenization and equilibration locally over the flux surfaces means that the thermal energy absorbed by the shielding cloud is immediately returned to the background plasma. We therefore do not need any further energy loss terms directly associated with the pellet ablation (assuming the shielding cloud is optically thick so that we may neglect radiative losses from it).

The fourth term in equation (4.8) describes a diffusive energy transport due to magnetic perturbations. When electrons follow stochastically perturbed magnetic field lines, their radial dynamics may be approximated as a diffusion process. In a tokamak geometry, a heuristic argument gives the Rechester-Rosenbluth form for the radial diffusion coefficient, $D = \pi q v_{\parallel} R_0 (\delta B/B)^2$ [62], for a particle traveling with a speed v_{\parallel} along the field lines. The factor $\pi q R_0$, where $q \approx 1$ is the safety factor, represents the parallel correlation length scale of the stochastic magnetic field perturbation, and $\delta B/B$ is the relative amplitude of the perturbation. The local electron heat diffusion coefficient D_W is calculated by integrating the diffusion coefficient D over a Maxwellian with the local temperature $T_{\rm M}$:

$$D_W = \frac{n_{\rm M}}{(\pi^{3/2} v_T^3 T_{\rm M})} \int \frac{m_e v^2}{2} \left(\frac{v^2}{v_T^2} - \frac{3}{2}\right) D(\boldsymbol{v}) \exp\left(-\frac{v^2}{v_T^2}\right) \mathrm{d}\boldsymbol{v}, \tag{4.13}$$

where $v_T = \sqrt{2T_M/m_e}$ is the electron thermal velocity. Lacking a self-consistent model for it, the evolution of $\delta B/B$ is prescribed in our model. When studying ITER-like scenarios, we typically set $\delta B/B$ to the order of 10^{-3} . This value results in a transport loss time scale a^2/D_W , with D_W evaluated at the initial temperature, of the same order of magnitude as the expected thermal quench time in ITER [42].

As a boundary condition for this diffusion, we assume $T_{\rm M} = 0$ at the plasma edge, but with D_W calculated at the temperature one grid point inside the edge to avoid having zero transport through the plasma edge.

The fifth term in equation (4.8) describes the collisional energy transfer from hot electrons to the Maxwellian electrons³. Here, we have $\Delta \dot{E}_{ee} = 4\pi n_{\rm M} r_0^2 \ln \Lambda_{ee} m_e c^4 / v$, where $r_0 = e^2 / (4\pi \varepsilon_0 m_e c^2)$ is the classical electron radius and $\ln \Lambda_{ee}$ is the energy dependent Coulomb logarithm for electron-electron collisions. The expression for $\ln \Lambda_{ee}$ was obtained in ref. [63] by matching the thermal Coulomb logarithm $\ln \Lambda_0$ with the relativistic Coulomb logarithm $\ln \Lambda_c = \ln \Lambda_0 + \ln \sqrt{m_e c^2 / T_{\rm M}}$ according to

$$\ln \Lambda_{ee} = \ln \Lambda_0 + \frac{1}{k} \ln \left(1 + \left(\frac{2(\gamma - 1)m_e c^2}{2T_{\rm M}} \right)^{k/2} \right), \tag{4.14}$$

where the interpolation parameter k was chosen to be k = 5. Finally, the last term in equation (4.8) accounts for the bremsstrahlung losses.

4.3 Electric field evolution

The rapid cooling of the plasma is accompanied by a rapid drop in the conductivity, resulting in the induction of an electric field, that later decays diffusively according to equation (2.42) derived in section 2.4.3. In the cylindrical approximation, equation (2.42) takes the form

$$\mu_0 \frac{\partial j_{||}}{\partial t} = \frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial E}{\partial r} \right), \tag{4.15}$$

where $j_{||}$ denotes the current density parallel to the field lines. The total current density $j_{||}$ is given by the sum of the Ohmic current density, carried by the Maxwellian electrons⁴, and the hot electron and runaway current densities. The Ohmic current density is calculated as $j_{\text{Ohm}} = E_{||}(\sigma_{\text{M}} - \sigma_{\text{fp}}) + \int_{0}^{p_{\text{hot}}} ev_{||}f \, d\boldsymbol{p}$, where σ_{M} is the conductivity given in equation (4.9). The first term compensates for the part of the conductivity not captured by the conductivity σ_{fp} resulting from the test-particle Fokker-Planck collision operator used in the kinetic equation described in section 4.4.1. An expression for σ_{fp} was determined by running numerous DREAM simulations with fixed parameters until the kinetically captured contribution to the Ohmic current was equilibrated, and then calculating σ_{fp} by dividing the Ohmic current thus obtained by the fixed value for $E_{||}$. This data was used to fit an expression for σ_{fp} according to

$$\sigma_{\rm fp} = \sigma_{\rm M} \left(1 + \frac{c_1}{c_2 + Z_{\rm eff}} \right), \tag{4.16}$$

³As this interaction term depends on the energy distribution, which is not resolved for the runaway population, only the hot population is accounted for in this interaction term.

⁴An exactly Maxwellian distribution does not carry a net current, so a small deviation from a Maxwellian is needed for the electrons we refer to as the Maxwellian population to carry the Ohmic current. The momentum scale of this deviation is however very small compared to the width of this close to Maxwellian distribution, and therefore this population can still be treated as a Maxwellian in other circumstances.

with $c_1 = -1.406$ and $c_2 = 1.888$. The hot electron current density is $j_{\text{hot}} = \int_{p_{\text{hot}} , and the runaway current density is <math>j_{\text{RE}} = ecn_{\text{RE}}$. The boundary condition for equation (4.15) at r = a is obtained by assuming the plasma to be surrounded by a perfectly conducting wall at r = b > a, where the electric field is set to zero. Matching the solution for r < a to the vacuum solution for a < r < b gives $E_{||}(a) = a \ln (a/b) \partial E_{||} / \partial r|_{r=a}$.

4.4 Runaway generation

The induction of a strong electric field results in an acceleration of electrons into the superthermal populations, eventually leading to the generation of a runaway current. In this work, we capture the seed runaway generation from the hot-tail, Dreicer, tritium decay, and Compton scattering mechanisms, as well as the subsequent runaway multiplication by the avalanche mechanism. The hot-tail mechanism is an intrinsically transient, kinetic effect resulting from the rapid temperature drop during the thermal quench and the finite equilibration time of the tail of the distribution function. Therefore, we invoke the kinetic equation described in section 4.4.1 during the thermal quench to model this effect. When the kinetic equation is invoked, it also captures the Dreicer mechanism. The runaway generation from tritium decay, Compton scattering, and avalanche are however not accounted for by the kinetic equation considered in this work, but are instead modeled as quasi-stationary sources feeding particles directly into the runaway population. After the thermal quench, when the hot-tail mechanism is no longer active, the Dreicer mechanism is also modeled in a similar fluid-like fashion. These source terms are described in section 4.4.2.

4.4.1 Kinetic equation

In order to capture the hot-tail mechanism, we consider the gyro-averaged kinetic equation with a linearised relativistic test particle Coulomb Fokker-Planck collision operator. The reference Maxwellian around which this collision operator is linearised is defined by the density $n_{\rm M}$ and temperature $T_{\rm M}$, introduced in section 4.1. As the kinetic equation is invoked here to study the comparatively small non-Maxwellian population, while the Maxwellian population is evolved by the fluid models described above, the field particle term in equation (2.26) is omitted (except that its effect on the Ohmic current is accounted for by the conductivity correction described in section 4.3). In addition to the test particle collision operator, we include a diffusive radial transport term with the Rechester-Rosenbluth form of the diffusion coefficient. The slowing down and deflection frequencies in the test particle operator are taken to be those accounting for the energy dependence of the Coulomb logarithm, as well as the effect of partial screening in collisions with partially ionized impurities, as derived in ref. [63]. Using the coordinates p and ξ , where p is the magnitude of the

momentum and ξ is the cosine of the pitch angle, this kinetic equation reads

$$\frac{\partial f}{\partial t} + eE\left(\frac{1}{p^2}\frac{\partial}{\partial p}\left[p^2\xi f\right] + \frac{1}{p}\frac{\partial}{\partial\xi}\left[\left(1-\xi^2\right)f\right]\right) = \frac{1}{p^2}\frac{\partial}{\partial p}\left[p^3\nu_s f\right] + \frac{\nu_D}{2}\frac{\partial}{\partial\xi}\left[\left(1-\xi^2\right)\frac{\partial f}{\partial\xi}\right] + \frac{1}{r}\frac{\partial}{\partial r}\left[rD\frac{\partial f}{\partial r}\right] + S\delta(\bar{\boldsymbol{p}}), \quad (4.17)$$

where we neglect the collisional energy-diffusion term, an assumption strictly valid in the superthermal limit. The strength S of the delta function source term at $\bar{\boldsymbol{p}} = \boldsymbol{p}/(m_e c) = 0$ is determined by the requirement that the total number of particles on the kinetic grid must satisfy quasi-neutrality, $\int f(\bar{\boldsymbol{p}}) d\bar{\boldsymbol{p}} = n_{\text{free}} - n_{\text{RE}}$.

The total slowing down and deflection frequencies, ν_s and ν_D , both contain contributions from electron-electron collisions and electron-ion collisions. The electronelectron contributions, denoted ν_s^{ee} and ν_D^{ee} , respectively, are given by [64]

$$\nu_{s}^{\text{ee}} = \nu_{c} \frac{\gamma^{2} \Psi_{1} - \Theta \Psi_{0} + (\Theta \gamma - 1) \bar{p} e^{-(\gamma - 1)/\Theta}}{\bar{p}^{3} e^{1/\Theta} K_{2}(1/\Theta)},$$

$$\nu_{D}^{\text{ee}} = \frac{\nu_{c}}{\gamma \bar{p}^{5} e^{1/\Theta} K_{2}(1/\Theta)} \left[(\bar{p}^{2} \gamma^{2} + \Theta^{2}) \Psi_{0} + \Theta(2\bar{p}^{4} - 1) \Psi_{1} + \gamma \Theta[1 + \Theta(2\bar{p}^{2} - 1)] \bar{p} e^{-(\gamma - 1)\Theta} \right],$$

$$\Psi_{n} = \int_{0}^{\bar{p}} (1 + s^{2})^{(n-1)/2} e^{-(\sqrt{1 + s^{2}} - 1)/\Theta} \, \mathrm{d}s,$$

$$\Theta = \frac{T_{\mathrm{M}}}{m_{e} c^{2}},$$

$$\nu_{c} = 4\pi \ln \Lambda_{ee} n_{\mathrm{M}} r_{0}^{2} c,$$
(4.18)

where K_2 is the second-order modified Bessel function of the second kind. The electron-ion contributions ν_s^{ei} and ν_D^{ei} , assumed to be against infinitely massive ions, with partial screening effects taken into account, are given by [63]

$$\nu_{s}^{ei} = 4\pi c r_{0}^{2} \frac{\gamma^{2}}{\bar{p}^{3}} \left(n_{\mathrm{M}} \ln \Lambda_{ee} + \sum_{ij} n_{ij} N_{e,ij} \left[\frac{1}{k} \ln(1 + h_{ij}^{k}) - \frac{\bar{p}^{2}}{\gamma^{2}} \right] \right),$$

$$\nu_{D}^{ei} = 4\pi c r_{0}^{2} \frac{\gamma}{\bar{p}^{3}} \sum_{ij} n_{ij} \left(\ln \Lambda_{ei} Z_{ij}^{2} + g_{ij}(\bar{p}) \right),$$

$$g_{ij} = \frac{2}{3} (Z_{\mathrm{tot},j}^{2} - Z_{ij}^{2}) \ln[1 + (\bar{a}_{ij}\bar{p})^{3/2}] - \frac{2}{3} N_{e,ij}^{2} \frac{(\bar{a}_{ij}\bar{p})^{3/2}}{1 + (\bar{a}_{ij}\bar{p})^{3/2}},$$

$$h_{ij} = \frac{m_{e}c^{2}}{\mathcal{I}_{ij}} \bar{p} \sqrt{\gamma - 1},$$

$$N_{e,ij} = Z_{\mathrm{tot},j} - Z_{ij},$$

$$\ln \Lambda_{ei} = \ln \Lambda_{0} + \frac{1}{k} \ln \left(1 + \left(\frac{2\bar{p}\sqrt{m_{e}c^{2}}}{\sqrt{2T_{\mathrm{M}}}} \right)^{k} \right),$$
(4.19)

where \mathcal{I}_{ij} is an ionic mean stopping power and \bar{a}_{ij} is an ion-specific screening length scale (normalised to the Bohr radius), both of which are tabulated in ref. [63].

Moreover, $\ln \Lambda_{ei}$ is the energy dependent Coulomb logarithm for electron-ion collisions, obtained in a similar way to $\ln \Lambda_{ee}$ described in section 4.2. The slowing down frequency $\nu_{\rm s}$ also contains an additional term $\nu_{\rm s}^{\rm br}$ accounting for the effect of radiation losses due to bremsstrahlung, based on a mean-force model [65], according to

$$\nu_{s}^{\rm br} = \frac{\gamma}{\bar{p}} \alpha r_{0}^{2} c \sum_{ij} n_{ij} Z_{\text{tot},j}^{2} \left[\frac{12\gamma^{2} + 4}{3\gamma\bar{p}} \ln(\gamma + \bar{p}) - \frac{8\gamma + 6\bar{p}}{3\gamma\bar{p}^{2}} \ln^{2}(\gamma + \bar{p}) - \frac{4}{3} + \frac{2F[2\bar{p}(\gamma + \bar{p})]}{\gamma\bar{p}} \right],$$

$$F(x) = \int_{0}^{x} \frac{\ln(1+y)}{y} dy,$$
(4.20)

where $\alpha \approx 1/137$ is the fine-structure constant. The effects of screening of partially ionized impurities on the bremsstrahlung emission have been ignored, so that the ions are only included in this expression through their total charge and density.

4.4.2 Fluid runaway sources

The time evolution of the runaway population is determined by the sum of the momentum space flux across p_{RE} from the hot population, the sources feeding particles directly into the runaway population, and the transport due to magnetic field perturbations:

$$\frac{\partial n_{\rm RE}}{\partial t} = F_p + \left(\frac{\partial n_{\rm RE}}{\partial t}\right)^{\rm Dreicer} + \left(\frac{\partial n_{\rm RE}}{\partial t}\right)^{\rm tritium} + \left(\frac{\partial n_{\rm RE}}{\partial t}\right)^{\gamma} + \left(\frac{\partial n_{\rm RE}}{\partial t}\right)^{\rm avalanche} + \frac{1}{r}\frac{\partial}{\partial r}\left[rD\frac{\partial n_{\rm RE}}{\partial r}\right],$$
(4.21)

where the flux F_p is given by $F_p = 2\pi p^2 \int_{p=p_{\rm RE}} (eE_{||}\xi - p\nu_s) f d\xi$, with the integral carried out along the upper momentum boundary $p_{\rm RE}$ of the domain representing the kinetic hot electrons. Note that this term is only active when the kinetic equation is invoked, and in that case it accounts for both the hot-tail and Dreicer mechanisms. When the kinetic equation is not invoked, the Dreicer mechanism is instead accounted for by the source term $(\partial n_{\rm RE}/\partial t)^{\rm Dreicer}$, calculated by a neural network trained on output from kinetic simulations [32].

The runaway seed produced by tritium decay is modeled as [16, 35]

$$\left(\frac{\partial n_{\rm RE}}{\partial t}\right)^{\rm tritium} = \ln\left(2\right)\frac{n_{\rm T}}{\tau_{\rm T}}f\left(W_{\rm crit}\right),\qquad(4.22)$$

where $n_{\rm T}$ is the tritium density, $\tau_{\rm T} \approx 4500$ days is the half-life of tritium. The fraction of the electrons created by tritium β^- decay above the critical runaway energy $W_{\rm crit}$ is given by $f(W_{\rm crit}) \approx 1 - (35/8)w^{3/2} + (21/4)w^{5/2} - (15/8)w^{7/2}$, with $w = W_{\rm crit}/Q$ and Q = 18.6 keV, corresponding to the maximum energy of the β^- electrons. The critical runaway energy $W_{\rm crit}$ is given by $W_{\rm crit} = m_{\rm e}c^2 \left(\sqrt{\bar{p}_{\star}^2 + 1} - 1\right)$,

in terms of the critical momentum for runaway acceleration (normalised to $m_e c$) with screening effects taken into account, implicitly given by

$$\bar{p}_{\star} = \frac{\sqrt[4]{\bar{\nu}_{\rm S}(\bar{p}_{\star})\bar{\nu}_{\rm D}(\bar{p}_{\star})}}{\sqrt{E_{||}/E_{\rm c}}},\tag{4.23}$$

as derived in Appendix A of ref. [19]. The collision frequencies $\bar{\nu}_{\rm s}$ and $\bar{\nu}_{\rm D}$ have here been normalised to ν_c , and $E_{\rm c} = \nu_c m_e c/e$ is the critical electric field for runaway generation introduced in section 2.3.3.

Runaway generation due to Compton scattering of gamma photons from the activated wall, taken from ref. [16], is based on the gamma photon energy spectrum

$$\Gamma_{\gamma}(E_{\gamma}) = \Gamma_{\gamma 0} \exp\left(-\exp\left(z\right) - z + 1\right),\tag{4.24}$$

with $z = [\ln (E_{\gamma} [\text{MeV}]) + 1.2] / 0.8$ and $\Gamma_{\gamma 0} = 4.44 \cdot 10^{17} \text{ m}^{-2} \text{s}^{-1} \text{MeV}^{-1}$, giving a total flux of $10^{18} \text{ m}^{-2} \text{s}^{-1}$, when integrated over the gamma photon energy E_{γ} . This spectrum was obtained using radiation transport calculations performed at several poloidal locations in ITER [16]. Using the expression for the total Compton cross-section [16]

$$\sigma(E_{\gamma}) = \frac{3\sigma_{\rm T}}{8} \left\{ \frac{x^2 - 2x - 2}{x^3} \ln \frac{1 + 2x}{1 + x(1 - \cos\theta_{\rm c})} + \frac{1}{2x} \left[\frac{1}{\left[1 + x(1 - \cos\theta_{\rm c})\right]^2} - \frac{1}{(1 + 2x)^2} \right] - \frac{1}{x^3} \left[1 - x - \frac{1 + 2x}{1 + x(1 - \cos\theta_{\rm c})} - x\cos\theta_{\rm c} \right] \right\},$$

$$(4.25)$$

the corresponding runaway generation rate can be evaluated as

$$\left(\frac{\partial n_{\rm RE}}{\partial t}\right)^{\gamma} = n_{\rm tot} \int \Gamma_{\gamma}(E_{\gamma}) \sigma(E_{\gamma}) dE_{\gamma}.$$
(4.26)

The critical Compton scattering angle θ_c for which an energy equal to W_{crit} is transferred to the electron from a photon with energy E_{γ} is given by

$$\cos \theta_{\rm c} = 1 - \frac{m_{\rm e}c^2}{E_{\gamma}} \frac{W_{\rm crit}/E_{\gamma}}{1 - (W_{\rm crit}/E_{\gamma})}.$$

The Thomson scattering cross section is $\sigma_{\rm T} = 8\pi/3[e^2/(4\pi\epsilon_0 m_{\rm e}c^2)]^2$, with $x = E_{\gamma}/(m_{\rm e}c^2)$, and $n_{\rm tot}$ is the total (free+bound) electron density. The dependence on $n_{\rm tot}$ rather than the free electron density reflects the fact that the gamma photon energies are much larger than the ionization energies of the ions and atoms in the plasma. Thus, the bound electrons may be scattered into the runaway momentum region essentially to the same extent as the free electrons.

The growth rate for the avalanche process is given by [20]

$$\left(\frac{\partial n_{\rm RE}}{\partial t}\right)^{\rm avalanche} = \frac{e n_{\rm RE}}{m_{\rm e} c \ln \Lambda_{\rm c}} \frac{n_{\rm tot}}{n_{\rm M}} \frac{E_{||} - E_{\rm c}^{\rm eff}}{\sqrt{4 + \bar{\nu}_{\rm S}(\bar{p}_{\star})\bar{\nu}_{\rm D}(\bar{p}_{\star})}},\tag{4.27}$$

where $E_{\rm c}^{\rm eff}$ is the effective critical electric field taking into account screening effects, as well as bremsstrahlung and synchrotron radiation losses due to the gyro-motion of the electrons, which were not accounted for in the expression introduced in section 2.3.3. The expression for $E_{\rm c}^{\rm eff}$ was derived in ref. [66]⁵. To obtain a well-behaved formula also for $E_{||} < E_{\rm c}^{\rm eff}$, which we use to approximately describe the runaway decay at near-critical electric fields (also in the absence of losses due to magnetic perturbations), we replace $p_{\star}(E_{||})$ by $p_{\star}(E_{\rm c}^{\rm eff})$ for $E_{||} < E_{\rm c}^{\rm eff}$. Finally, the last term in equation (4.21) accounts for the transport of runaway electrons due to magnetic field perturbations, with the Rechester-Rosenbluth diffusion coefficient D evaluated for $v_{||} = c$.

4.5 Numerical implementation and settings

Having described our model from a physics point of view, we here give a brief overview of the numerical tool DREAM (Disruption Runaway Electron Avoidance Model) [25], and typical settings used in the disruption studies presented in chapter 5. Based on an input specifying the injection parameters and pre-disruption plasma conditions, DREAM is capable of self-consistently calculating the time evolution of the background plasma properties, the electron momentum distribution and the runaway current during a mitigated tokamak disruption. Specifically, the user first defines plasma geometry, determined by the plasma minor radius a, the minor radius of the tokamak wall b, and the tokamak major radius R_0 . The user also defines the SPI settings, defined by the number N_{inj} of injected particles, the number of shards $N_{\rm s}$ into which the pellet is shattered, the particle fraction ϕ_i of the pellet consisting of species j, the average shard speed $\langle v_{\rm p} \rangle$ and corresponding distribution width $\Delta v_{\rm p}$, and the divergence angle α_p . Finally, the user provides initial profiles of the current density j_{\parallel} , temperature $T_{\rm M}$ and ion densities n_{ij} . Based on this input, the initial conditions for the pellet shard parameters are chosen randomly; with a shard size distribution given by equation (4.1) and uniform speed and angular distributions. The initial profile for j_{\parallel} is assumed to be purely Ohmic, which also imposes an initial condition on the electric field, as $E_{||} = j_{Ohm}/\sigma_M = j_{||}/\sigma_M$. The input profiles for T_M and n_{ij} are used directly as initial conditions, with n_{ij} typically describing a fully ionized deuterium or deuterium-tritium plasma. The initial momentum distribution of the electrons is chosen to be a Maxwellian at the local initial temperature.

From these input parameters and associated initial conditions, DREAM calculates the subsequent time evolution of the pellet and plasma parameters, the model of which can be summarised as follows. The pellet shards follow straight lines as described by equation (4.2), and the evolution of the shard sizes is governed by the ablation rate (equation (4.3)). The resulting density increase is given by equation (4.6), and the resulting cooling is calculated by equation (4.8). Once the pellet material is deposited, the evolution of n_{ij} is governed by ionization and recombination according to equation (4.10). The electric field evolution resulting from the rapid change in the conductivity is given by the induction/diffusion process described by equation (4.15).

⁵A numerical implementation of E_{c}^{eff} is available at https://github.com/hesslow/Eceff

The momentum distribution of the electrons is resolved in the range 0 ,and is evolved using the kinetic equation (4.17) (referred to as the "fully kinetic"mode in ref. [25]). The distribution thus calculated accounts for the current carriedby the (essentially) Maxwellian and hot electron population, complemented with a $contribution <math>E_{\parallel}(\sigma_{\rm M} - \sigma_{\rm fp})$, accounting for the field particle term not being included in equation (4.17), as described in section 4.3. Finally, electrons passing the upper momentum boundary, $p_{\rm RE}$, are added to the fluid-like runaway population, together with the contributions from the additional runaway sources summarised by equation (4.21), which conclude the equation system. The self-consistent solution of this equation system, performed with DREAM, is now briefly described below.

The equation system is discretized using a finite volume method (see e.g. ref. [67]) for the discretization of the momentum and configuration space, and an Euler backward scheme for the time discretization. The momentum space is divided into $N_{\rm p} \times N_{\xi}$ grid cells, the configuration space into $N_{\rm r}$ grid cells, and the time is discretized with a time step Δt . The momentum scale of the features of the distribution function are typically much smaller at low momenta compared to high momenta towards the end of the temperature drop, when the momentum scale of the Naxwellian population becomes rather narrow. In order to properly resolve the low energy region throughout the whole disruption without over-resolving the high energy region, the momentum grid is therefore divided into two parts with different resolutions, separated at a momentum $\bar{p}_{\rm sep}^{6}$. The resulting momentum grid is discretized by $N_{\rm p1}$ uniformly spaced grid cells between $0 < \bar{p} < \bar{p}_{\rm sep}$ and $N_{\rm p2}$ uniformly spaced grid cells between $\bar{p}_{\rm sep} < \bar{p} < \bar{p}_{\rm RE}$.

Approximating the derivatives with central differences in agreement with the discretization, the differential equations included in the model are translated into an algebraic equation system for the cell averages of the various quantities at the next time step. The evolved quantities are thus computed in the center of the cells, while the fluxes between adjacent cells are calculated on the cell surfaces, ensuring that the flux into a grid cell exactly equals the flux out of adjacent grid cells. This guarantees conservation of the integrals of the various quantities within machine precision, in the absence of sources and edge losses, thus satisfying the physical conservation laws of e.g. particle number and energy. The interpolation from the center of the cells to the cell boundaries, needed for the calculation of the fluxes, is performed in a way as to preserve positivity of the evolved quantities.

The implicit solution for the evolved quantities in the next time step is obtained via Newton iteration. Letting \mathbf{X}^i denote a vector containing the cell averages of every quantity in every cell at time step *i*, the equation system to be solved can be written in the form $\mathbf{F}(\mathbf{X}^{i+1}, \Delta t) = \mathbf{0}$. The Newton iteration scheme can then be expressed as

$$\boldsymbol{X}_{j+1}^{i+1} = \boldsymbol{X}_{j}^{i+1} - J^{-1}(\boldsymbol{X}_{j}^{i+1}, \Delta t) \boldsymbol{F}(\boldsymbol{X}_{j}^{i+1}, \Delta t), \qquad (4.28)$$

where the subscript on X denotes the iteration number and $J^{-1}(X_j^{i+1}, \Delta t)$ is the inverse of the Jacobian of $F(X_j^{i+1}, \Delta t)$. Here we have also explicitly stated the dependence of J^{-1} and F on the time step Δt , which enters through the discretized

 $^{^6\}mathrm{Not}$ to be confused with the momentum \bar{p}_hot separating the Maxwellian and hot electron population.

time derivatives. The Jacobian is constructed based on analytical derivatives, although some are approximated or even neglected for simplicity and/or limiting the complexity of the construction and inversion of the Jacobian. The iteration starts at the quantity vector in the previous time-step, i.e. $X_0^{i+1} = X^i$, and continues until a tolerance specified by the user is satisfied separately for every evolved quantity. Specifically, the convergence condition for the part $X_{n,j+1}^{i+1}$ corresponding to quantity n is

$$||\boldsymbol{X}_{n,j+1}^{i+1} - \boldsymbol{X}_{n,j}^{i+1}||_{2} \le \epsilon_{n}^{\text{abs}} + \epsilon_{n}^{\text{rel}}||\boldsymbol{X}_{n,j+1}^{i+1}||_{2},$$
(4.29)

where $|| \cdot ||_2$ denotes the 2-norm and the absolute and relative tolerances ϵ_n^{abs} and ϵ_n^{rel} are specified by the user.

A detailed description and demonstration of the core functionality of DREAM has recently been compiled and submitted for publication elsewhere [25]. In this work, the functionality of DREAM has been extended to include the SPI model described in section 4.1. A non-trivial design choice related to the extension implemented in this work concerns the singularity of the ablation rate in equation (4.3) at $r_{p,k} = 0$. To circumvent this issue, we remove the singularity by a change of variable according to $Y_{p,k} = r_{p,k}^{5/3}$, such that the ablation rate can be expressed as

$$\dot{Y}_{\mathrm{p},k} = -\lambda(X) \frac{5}{3r_0^{4/3}} \left(\frac{q_{\mathrm{in}}}{q_0}\right)^{1/3} \left(\frac{\mathcal{E}_{\mathrm{in}}}{\mathcal{E}_0}\right)^{7/6} \frac{1}{4\pi\rho_{\mathrm{dens}}},\tag{4.30}$$

for $\dot{Y}_{\mathrm{p},k} > 0$. The discretized density source terms, corresponding to equation (4.6), are related to the change of $\dot{Y}_{\mathrm{p},k}$ in a particle conserving manner, assuming a delta-like deposition kernel whose discretisation is described in Appendix A. The density source terms contain contributions from all shards for which $Y_{\mathrm{p},k} > 0$ in the beginning of the time step. However, small negative values of $Y_{\mathrm{p},k} > 0$ and the corresponding density source term are allowed during an individual time step, to avoid convergence issues for the Newton iteration process related to a non-smooth Jacobian at $Y_{\mathrm{p},k} = 0$.

In addition to the input related to the physical state of the system, DREAM also requires specification of the numerical parameters regarding the discretization and Newton iteration process. The most computationally efficient parameters vary from case to case, and also between the various parts of a disruption, but here we give a set of typical numerical parameters for the simulations shown in chapter 5. By default, the relative tolerance $\epsilon_n^{\text{rel}} = 10^{-6}$, and the absolute tolerance is disabled, i.e. $\epsilon_n^{\text{abs}} = 0$. Furthermore, we typically set the time step to the order of 1-10 µs, choosing a time step in the shorter end of this range during the part of the thermal quench when the majority of the thermal energy loss takes place and in the longer end of this range for the subsequent current quench and runaway plateau.

In order to resolve the cold electron bulk at the end of the thermal quench, a momentum resolution of the bulk corresponding to a sub-eV increment in energy is required. Therefore, we discretize the momentum grid below $\bar{p}_{sep} = 0.07$ with a comparatively high resolution of $N_{p1} = 70$ grid cells uniformly spaced between $\bar{p} = 0$ and $\bar{p} = \bar{p}_{sep}$. The features of the hot tail, which can form at higher momenta, have a substantially larger momentum scale, and therefore do not require as high resolution. For this part of the momentum grid, with $\bar{p}_{sep} < \bar{p} < \bar{p}_{RE} = 3$, we use $N_{p2} = 50$ uniformly spaced grid cells in \bar{p} . In all of momentum space, the pitch angle

variable ξ is discretized by $N_{\xi} = 5$ grid cells uniformly spaced between -1 and 1. Finally, we discretize the radial dimension with $N_r = 11$ grid cells. This resolution was found to be sufficient to resolve the features of interest in this work, while still enabling extensive scans in the SPI parameters of interest, even with the somewhat computationally expensive thermal quench simulations with the kinetic part of the model invoked.

4. Disruption model

5

Simulations of two-stage SPI injections

Having established the SPI and disruption model in chapter 4, we are now ready to address the main goal of this work, which is to evaluate the disruption mitigation performance in an ITER-like plasma of the two-stage deuterium-neon injection scheme introduced in section 3.4.2.2. With this scheme, the plasma is first cooled by dilution down to the 100 - 1000 eV range by a deuterium injection, without significantly affecting the thermal energy density (and hence the pressure). A few ms later, neon is injected, with the aim to radiate away most of the thermal energy. A major advantage of the temporal separation between the injections is that the intermediate equilibration of the electron momentum distribution has the potential to significantly reduce the hot-tail runaway generation. Another advantage is that the radiative loss of thermal energy becomes larger compared to the transport losses if the thermal quench is triggered from a lower temperature. This is advantageous because it would help to spread the heat load on the surrounding structures over a wider area. These aspects are therefore particularly emphasised in our investigation.

The rest of this thesis is structured as follows: in section 5.1, we investigate the effects of the SPI engineering parameters defined in section 4.1 on the deposition profile of the first deuterium injection. The goal here is to find parameters where the density increase covers the whole plasma, while still minimizing the amount of pellet material passing through the core without being ablated. Section 5.2 contains a similar investigation for the secondary neon injection following a representative example of a primary deuterium injection. The dynamics of a typical two-stage injection are also compared to the dynamics of a single stage injection with similar injected quantities of deuterium and neon in section 5.2.1. Having found suitable parameters for efficient use of the pellet material, the radiated fraction of the thermal energy as a function of the number of injected deuterium and neon particles is investigated in section 5.2.2. Next, we invoke the kinetic equation from section 4.4.1to study the reduction of the hot-tail runaway generation during a two-stage SPI compared to a single stage SPI in section 5.2.3. The subsequent current evolution during the current quench is studied in section 5.3, including the hot-tail seed as well as the other primary sources and the avalanche mechanism described in section 4.4.2. The implications and validity of these results are discussed in chapter 6.

5.1 First stage deuterium injections

We start with addressing the problem of finding parameters which give an efficient density increase in the first-stage deuterium injection. The parameters to be set, defined in section 4.1, are the number of injected deuterium and neon particles, $N_{\rm inj,D}$ and $N_{\rm inj,Ne}$, the number of shards $N_{\rm s,D}$ and $N_{\rm s,Ne}$ the pellets are shattered into, the pellet speeds and variances $\langle v_{\rm p,D} \rangle$, $\Delta v_{\rm p,D}$, $\langle v_{\rm p,Ne} \rangle$ and $\Delta v_{\rm p,Ne}$, and the divergence angles $\Delta \alpha_{\rm p,D}$ and $\Delta \alpha_{\rm p,Ne}$ of the injected shards.

In order to achieve the desired dilution of the initially hot plasma, from an initial temperature in the 10 keV range down to the 100-1000 eV range, an electron density increase of the order of 10-100 times the initial density is required. The corresponding number of assimilated particles is of the order of $10^{24} - 10^{25}$. Moreover, the requirement of a fast reaction time of the disruption mitigation system favours fast pellet speeds. A fast pellet speed also makes it possible to achieve core penetration with a smaller amount of injected particles. In that way, using fast pellets extends the available parameter space of post-injection densities compared to a slower injection. For large injections, the assimilation might have to be enhanced to avoid shards passing through the plasma without ablating. This can be done by shattering the pellet into more shards, as discussed in section 3.4.1.3.

In this chapter, we search for suitable injection parameters for ITER-like plasmas with initial temperature profile $T_{\rm M}(r) = T_0(1 - 0.99(r/a)^2)$, central temperature $T_0 = 20$ keV, and initial density $n_{\rm M} = 10^{20}$ m⁻³ at all radii. The profiles used here have previously been used to study massive material injections in ITER-like scenarios [19], assuming flat deposition profiles. For the initial ion composition, we consider both a pure deuterium plasma and an even mix of deuterium and tritium, the latter corresponding to the nuclear operation phase. The plasma minor radius is a = 2 m, the radius of the conducting wall (used to set a boundary condition on the electric field, see 4.3) is b = 2.15 m and the major radius of the tokamak is $R_0 = 6.2$ m. The initial plasma current is $I_{\rm p} = 15$ MA, with a radial profile given by $j_{\parallel}(r) = j_0(1 - (r/a)^2)^{0.41}$, with $j_0 = 1.69$ MA/m². While the pellet shards are still present in the plasma, the number of superthermal electrons was found to be too small to have a significant impact on any of the plasma parameters relevant for the ablation of the shards. The superthermal electrons are therefore ignored in the simulations shown before section 5.2.3.

As explained in the beginning of this section, it is beneficial to have fast pellet speeds. We therefore consider the fastest expected injection speeds possible. We recall from section 3.4.1.3 that these are $\langle v_{\rm p,D} \rangle = 800$ m/s for deuterium pellets and $\langle v_{\rm p,Ne} \rangle = 200$ m/s for neon pellets. For the distribution of pellet speeds, we assume $\Delta v_{\rm p} = 0.2 \langle v_{\rm p} \rangle$ for both deuterium and neon pellets. Moreover, we fix the divergence angle $\Delta \alpha_{\rm p} = 20^{\circ}$ for both deuterium and neon pellets. This parameter primarily affects how many shards pass through the innermost flux surfaces, and an increased divergence can therefore shift the deposited density profile from the core further out in the plasma. While this property might be of interest when fine-tuning the density profile, we are able to achieve satisfactory density profiles with our fixed value of $\Delta \alpha_{\rm p}$. Fixing this value is further motivated by the remark in ref. [19], that the current quench dynamics is rather insensitive to the details of the density profile for a given number of injected particles, as long as core penetration is achieved. We also fix the position at which the pellets are shattered to the tokamak wall on the horizontal mid-plane, i.e. at position $(x_0, y_0) = (b, 0)$.

The injection parameters left to investigate are the number of injected deuterium and neon particles $N_{inj,D}$ and $N_{inj,Ne}$, and the number of shards $N_{s,D}$ and $N_{s,Ne}$ into which the pellets are shattered. In order to make the most efficient use of the pellet material, the number of shards for a given number of injected particles should be chosen to achieve core penetration without leaving any unablated pellet material. To this end, we perform a scan of the assimilation rate, i.e. the fraction of the pellet material that is deposited in the plasma, as a function of the number of injected particles and the number of shards. Consistently with the assumption of no correlations between the injection parameters mentioned in section 4.1, these parameters are varied independently of each other.

The deuterium assimilation rate as a function of $N_{\rm inj,D}$ and $N_{\rm s,D}$ is shown in figure 5.1 a). Core penetration is achieved in the region of this parameter space to the right of the green line. More specifically, here we define core penetration as obtaining a density increase in the core of 10% of the largest density increase in the plasma. For a large number of shards, the green line is expected to mostly follow the contour marking 100% assimilation rate. The slight discrepancy seen in figure 5.1 a) is explained by the finite radial resolution. For a low number of shards, however, the probability that enough shards will pass through the small innermost flux surfaces decreases, making core penetration less likely. This can affect the assimilation rate significantly, and is the reason behind the horizontal part of the green line at lower $N_{\rm s,D}$. This study now provides us with an optimal way of choosing $N_{\rm s,D}$ for a given $N_{\rm inj,D}$ (within the model at hand), by choosing $N_{\rm s,D}$ close to the 100% assimilation contour and to the right of the green line in figure 5.1 a). However, to make sure that a sufficient core density is achieved, we include a safety margin and choose $N_{\rm s,D}$ along the 97% contour.

Figure 5.1 b-d) show the profiles of the deuterium density, temperature and thermal energy density after the first injection stage. Specifically, the profiles are shown at 1 ms after the time $t_{\text{pass},D}$ when all deuterium shards have passed through the center (or would have if they were not completely ablated). All three panels show profiles for $N_{\rm inj,D} = 1.58 \cdot 10^{24}$ and a varying number of shards. We see in panel b) that for the smaller numbers of shards, the density increase peaks close to the core. This behaviour is expected, as the initial temperature, and hence the ablation rate, is the highest in the core. The size of a flux surface is also proportional to its radius r, so that more particles have to be deposited to achieve a given density increase further out in the plasma. Note, however, that for the innermost flux surfaces, the covered angle is smaller than the divergence of the plume of shards. The number of shards passing through these flux surfaces therefore decreases with decreasing r, moving the density peak away from the core. Moreover, as the number of shards increases, so does the ablation rate, as discussed in the beginning of this section. This makes the shards ablate more of their material earlier along their trajectories, moving the maximum density even further outwards. If the average shard size becomes small enough, such as in the $N_{s,D} = 3000$ case in figure 5.1, all shards burn out before they reach the core.

In panel c), we see that the post-injection temperature is of the order of a few hundred eV, as desired, in the regions that were effectively reached by the shards. However, in some cases there are areas with little or no deposition of pellet material, and in these areas the temperature remains higher. The latter is most clearly seen on the temperature profile for $N_{\rm s,D} = 3000$, where the temperature remains in the initial ~ 20 keV range in the innermost part of the plasma, which is not reached by any pellet shards.

Finally, we also note in panel d) that the energy density, and hence the pressure, is only slightly affected by the pellets. The reason for this is that the pure deuterium pellets do not cause significant radiative losses, but only dilute the thermal energy in the plasma. Keeping the pressure perturbation small during this first injection stage is essential, because it favours avoiding an unwanted acceleration of the growth of MHD instabilities. As described in the introduction of this chapter, thermal energy is instead intended to be dissipated radiatively by the neon in the second injection stage, which is the subject of the next section.



Figure 5.1: a) Assimilation rate as a function of the number of injected deuterium particles $N_{\rm inj,D}$ and number of shards $N_{\rm s,D}$. The green line marks the region of the parameter space where core penetration is achieved, defining core penetration as obtaining a density increase in the core of 10% of the largest density increase in the plasma. The other panels show b) deuterium density, c) temperature and d) thermal energy density 1 ms after $t_{\rm pass,D}$, resulting from pure deuterium SPIs with $N_{\rm inj,D} = 1.58 \cdot 10^{24}$ and varying numbers of shards (colour-code given in panel b).

5.2 Second stage neon injections

Having studied the outcome of the first-stage deuterium injection for a range of controlling parameters, we now perform a similar analysis of the outcome of the secondary neon injection. As the purpose of this neon injection is to dissipate energy through radiation rather than cooling by dilution, the desired neon density is orders of magnitude smaller than the injected deuterium density. Informed by the range of neon densities found in ref. [19] to give the most favourable current quench times and runaway currents, we choose to scan logarithmically in the range $10^{22} < N_{\rm inj,Ne} < 10^{24}$. As the neon pellets considered here are smaller than the deuterium pellets studied in section 5.1, we also consider somewhat smaller numbers of neon shards. The delay of the neon injection is chosen so that the neon shards enter the plasma edge at $t = t_{\rm pass,D} = 3.4$ ms after the shattering of the deuterium pellet.

Figure 5.2 a) shows the assimilation rate as a function of $N_{\rm inj,Ne}$ and $N_{\rm s,Ne}$, for a neon injection following a first stage deuterium injection with $N_{\rm inj,D} = 10^{24}$ and $N_{\rm s,D} = 66$. The number of deuterium shards was chosen along the 97% assimilation contour in figure 5.1. As the initial temperature is now only a few hundred eV, only the smallest neon pellets considered have the majority of their content ablated. For the largest neon pellets, the assimilation rate is only a few percent. The assimilation rate improves when the number of shards increases, but we note that this trend slows down at $N_{\rm s,Ne} \sim 50$. This is due to the first shards arriving at a given flux surface causing enough cooling to reduce the ablation rate for the later shards to a negligible level, resulting in a saturation of the assimilation rate.

The other panels in figure 5.2 show the neon density, temperature and thermal energy density 1 ms after $t_{\text{pass,Ne}} = 17.5$ ms (defined similarly to $t_{\text{pass,D}}$ but for the neon shards). The injection parameters shown are $N_{\rm inj,Ne} = 10^{23}$ and a varying number of shards. Similarly to the deuterium density profiles, the neon density profiles peak slightly outside the core, due to the dependence of the flux surface size and number of particles passing through the flux surfaces on their radius. The neon densities are, however, approximately 1-2 orders of magnitude lower than the deuterium densities. The final temperature has now decayed to the 5 eV range, where further decay is prevented by the Ohmic heating from the plasma current. Finally, we note that the thermal energy density is essentially completely dissipated by the neon radiation. The associated pressure gradient developing as the shards move through the plasma is expected to result in a rapid growth of MHD instabilities, and therefore a substantial magnetic perturbation should be accounted for when the neon shards enter the plasma. However, as the temperature is already relatively low when the neon shards enter the plasma, the effect of this magnetic perturbation will be found to be quite moderate when investigated in section 5.2.1 and 5.2.2.


Figure 5.2: a) Assimilation rate for the secondary neon injection as a function of $N_{\rm inj,Ne}$ and $N_{\rm s,Ne}$, following a deuterium SPI with $N_{\rm inj,D} = 10^{24}$ and $N_{\rm s,D} = 66$. The number of deuterium shards was chosen along the contour marking 97% assimilation in figure 5.1 a). The remaining panels show the neon density (b), temperature (c) and thermal energy density (d) 1 ms after $t_{\rm pass,Ne}$, for $N_{\rm inj,Ne} = 10^{23}$ and a varying number of shards (colour-code given in panel b).

5.2.1 Single vs two-stage injection dynamics

We continue our analysis by taking a closer look at the spatio-temporal evolution of the most relevant plasma parameters during a representative two-stage injection. As a reference, we construct a single-stage injection of a mixed neon-deuterium pellet containing the same number of injected neon and deuterium particles. For the two-stage injection, we use $N_{\rm inj,D} = 2 \cdot 10^{24}$, $N_{\rm s,D} = 1742$, $N_{\rm inj,Ne} = 10^{23}$ and $N_{\rm s,Ne} = 50$. The number of shards was again chosen along the 97% contour in figure 5.1 a). For the single-stage injection, we use a pellet consisting of 95% deuterium and 5% neon, with a total of $2 \cdot 10^{24}$ atoms shattered into 1742 shards. The larger mass of the mixed pellet lowers the maximum achievable pellet speed (see section 3.4.1.3), and we therefore lower $\langle v_{\rm p} \rangle$ to 600 m/s for the mixed neon-deuterium pellet. Moreover, we model the effect of magnetic perturbations with an amplitude of $\delta B/B = 0.001$ (assumed to be constant in space and follow a step function in time) when the neon shards (or mixed deuterium-neon shards) enter the plasma. This is done in order to emulate the stochastisation of the magnetic field resulting from the perturbations of the plasma caused by the pellet, as discussed in section 4.2. The magnetic perturbation is assumed to be active during the entire part of the disruption studied in this section.

The spatio-temporal evolution of the neon and deuterium density, thermal energy, temperature, electric field and current density, for the two-stage injection, are shown in figure 5.3. The two solid green lines mark the trace of the fastest and slowest deuterium shards (for shards with $\alpha_p = 0$), and the dashed green lines mark the corresponding trace for the neon shards. In the upper panels of figure 5.3, we see a clear increase in the deuterium density and neon density between the corresponding pair of green lines. However, the increase in the neon density saturates quite closely to the first dashed green line, especially in the inner part of the plasma. The reason for this is that the neon deposited from the first arriving shards radiatively cool the plasma enough to impede the ablation of the later shards.

Looking at figure 5.3 c-d), we again see that the thermal energy density is only slightly affected by the deuterium injection, while the temperature falls to a few hundred eV. When the neon shards enter the plasma, the thermal energy density is dissipated over a millisecond time scale or faster in the parts of the plasma that have been reached by the neon shards. Consequently, the temperature drops to a few eV on a similar time scale. The onset of the magnetic perturbations causes some diffusion of the thermal energy density in the parts of the plasma that have not yet been reached by the neon shards. However, the resulting contribution to the dissipation rate of the thermal energy is still quite moderate, in total dissipating 6% of the initial thermal energy. Notably, this value is smaller than the 10% limit mentioned in section 3.2, and much smaller than the radiative losses that dissipate almost all of the remaining thermal energy. Although it should be emphasised that our model for the thermal energy is rather crude, assuming an immediate onset of a prescribed magnetic perturbation that then remains constant in time and space, it can still be of interest for comparing the calculated transported losses in different scenarios.

In figure 5.3 e-f), we see the strong increase in the electric field in connection with the second temperature drop, and the start of the subsequent current decay. We also see a radial spike in the current profile moving inwards along with the neon shards. This spike is caused by the diffusion of the electric field, induced where the plasma has been cooled by the neon shards, into the still hotter region that the neon shards have not yet reached. In the hotter region, the conductivity is still high enough that even a relatively modest increase of the electric field can cause a significant increase in the (Ohmic) current density. As the neon shards move further into the plasma, so does the location of the conductivity drop, resulting in a rather fast resistive reduction of the radial current spike. This process is then repeated, causing the radial current spike to move inwards until both the neon shards and the current spike reach the core. When the temperature makes a further drop to around 1 eV in the outer part of the plasma, a similar set of spikes in the electric field are seen. This temperature drop occurs when the Ohmic current density is sufficiently low so that the Ohmic heating can no longer sustain a temperature in the few eV range. The conductivity in the nearby plasma is however too low for another set of

spikes in the current density to form.

When the neon shards have reached the core, in most cases, the current starts to decay in all parts of the plasma. As we will see in section 5.3, however, in some cases the Ohmic heating, amplified by the radial current spike, can cause parts of the plasma to re-heat. The re-heating is accompanied by an increase in the conductivity, and as the electric field diffuses into the re-heated regions the current density can initially increase locally. The decay of the total current is also comparatively slow, with a time scale of the order of seconds.



Figure 5.3: Spatio-temporal evolution of the deuterium density (a), neon density (b), thermal energy density (c), temperature (d), electric field (e) and current density (f) during a two-stage SPI injection with parameters $N_{\rm s,D} = 1724$, $N_{\rm inj,D} = 2 \cdot 10^{24}$, $N_{\rm s,Ne} = 50$, and $N_{\rm inj,Ne} = 1 \cdot 10^{23}$. The mean speed of the deuterium shards is $\langle v_{\rm p,D} \rangle = 800$ m/s and the mean speed of the neon shards is $\langle v_{\rm p,Ne} \rangle = 200$ m/s. A diffusive heat transport with $\delta B/B = 0.001$, constant in space and time, is activated once the neon shards enter the plasma.

Turning now to the single-stage injection, one of the most noticeable differences compared to the two-stage injection is that the neon and deuterium densities now increase simultaneously, as seen in figures 5.4 a) and b). The neon and the deuterium are now delivered by the same pellet, and therefore their radial density profiles are proportional. As the diffusion of the thermal energy now starts while the plasma still has a temperature of several keV, it has a much more substantial effect on the thermal energy density and the temperature compared to the two-stage injection, as seen in panels c) and d).

The diffusive decay of the temperature in the inner parts of the plasma before the arrival of the shards reduces the ablation there. The resulting density profiles therefore have the their maximum moved outwards compared to the density profiles of the two-stage injection. The fast dilution cooling by the deposited deuterium however still reduces the transport coefficient for the thermal energy in the regions passed by the shards. This gives rise to a ~ 0.5 ms long plateau-like phase in the local temperature evolution, reminiscent of the temperature plateau between the injections in the two-stage case. As a result, the fraction of the thermal energy lost by transport is limited to 24%. This is however substantially larger than the 6% obtained for the two-stage injection. This reduced transport also causes thermal energy to pile up in the cold region, giving rise to the off-center maximum aligned with the shard trajectories in the r - t diagram in figure 5.4 c). The evolution of the electric field and current density essentially follow the temperature evolution in a similar way as seen in the two-stage injection.



Figure 5.4: Spatio-temporal evolution of the deuterium density (a), neon density (b), thermal energy density (c), temperature (d), electric field (e) and current density (f) during a single stage mixed neon-deuterium SPI, with parameters chosen to give similar assimilated densities as the two-stage injection shown in figure 5.3. The pellet consists of a mixture of 5% neon and 95% deuterium, with $2 \cdot 10^{24}$ particles shattered into 1742 shards. A diffusive heat transport with $\delta B/B = 0.001$, constant in space and time, is activated once the mixed neon-deuterium shards enter the plasma.

5.2.2 Radiative vs transport losses

The comparison between the single-stage and two-stage injection in section 5.2.1 confirms one of the expected advantages of the two-stage injection scheme: that the transport losses of thermal energy are significantly lower than for the single-stage injection. In this section, we make a further study of the fraction of the thermal energy lost by transport during two-stage injections with different numbers of injected neon and deuterium particles.

For this analysis, we again use a deuterium pellet speed of $\langle v_{\rm p,D} \rangle = 800$ m/s. The number of shards the deuterium pellet is shattered into, for a given number of injected particles, is chosen along the 97% assimilation contour of figure 5.1 a), when possible. To illustrate the consequences of not achieving core penetration, we also include simulations of smaller deuterium pellets which are fully ablated before they reach the core regardless of the number of shards they are shattered into. For these pellets, we chose $N_{\rm s,D} = 10$. For the neon pellets, the relatively low temperature of the diluted plasma when they are injected makes full assimilation difficult for the larger pellets, as seen in figure 5.2. Moreover, it can be seen in figure 5.2 that the increasing trend in the assimilation rate when increasing the number of shards appears to slow down for shard numbers $\gtrsim 50$. Therefore, we fix the number of neon shards to $N_{\rm s,Ne} = 50$.

Figure 5.5 shows the fraction of the initial thermal energy that is lost by transport during a two stage injection as a function of the number of injected deuterium and neon particles. Here we see that the transported fraction can be significantly decreased by increasing the deuterium and neon content. The dependence on $N_{\rm ini,Ne}$ is however weaker than the dependence on $N_{inj,D}$, primarily due to the decrease in the assimilated fraction as $N_{inj,Ne}$ increases. The neon content primarily increases the radiative losses, which are further enhanced by the increase in the electron density due to the deuterium injection, as may be anticipated from the second term in equation (4.8). As discussed in section 5.2.1, the deuterium content also limits the transport by lowering the temperature before the onset of the magnetic perturbation. Again, we note that our model for the thermal energy transport is not detailed enough to give reliable quantitative predictions, but figure 5.5 may still be indicative of the trends regarding the reduction of the transported losses during a two-stage deuterium-neon SPI. It is also noteworthy that simulated transported losses below the limit of 10% of the initial thermal energy, marked by the dashed grey line in figure 5.5, are achievable within a realistic range of injection parameters.

Finally, we note that lack of core penetration does not cause a dramatic difference in the transported fraction. One reason for this is that the neon shards can reach the core even if the deuterium shards did not. Another reason is that the thermal energy in the core has to pass the outer regions of the plasma before it can be lost due to transport, and can therefore be radiated away in the regions that have been reached by the deuterium shards.



Figure 5.5: Fraction of the thermal energy lost by transport due to the magnetic field perturbations, as a function of the number of injected deuterium and neon particles. The parameters for the injections were $N_{\rm s,Ne} = 50$ for the neon shards, and the number of deuterium shards was chosen along the 97% assimilation contour in figure 5.1, when possible. This region is marked by the dotted grey line. To the left of the grey line, the deuterium pellets were shattered into 10 shards. A diffusive heat transport with $\delta B/B = 0.001$, constant in space and time, is activated once the neon shards enter the plasma. A dashed grey line marks the target of a transported fraction lower than 10%.

5.2.3 Hot-tail generation

We now turn to the other major expected advantage of the two-stage injection scheme compared to single-stage injections, which is the reduction of the hot-tail generation due to the intermediate equilibration of the electron momentum distribution. For this analysis, the hot-tail generation is captured by invoking the kinetic equation described in section 4.4.1. The fluid-like sources described in section 4.4.2 are still turned off, in order to distinguish the hot-tail seed generation (the Dreicer mechanism is also inherently accounted for by the kinetic equation in section 4.4.1, but gives here only a negligible contribution). In our model, this seed generation is given by the flux across the upper momentum boundary $p_{\rm RE}$ of the kinetic grid, defined in section 4.2. The total runaway current resulting from the hot-tail seeds found here is studied later in section 5.3, including all runaway generation mechanisms described in section 4.4.2.

As a reference case, we start by studying the hot-tail generation of the singlestage injection studied previously in figure 5.4. The corresponding spatio-temporal evolution of the runaway current density is shown in figure 5.6. We consider both the case where the superthermal electrons (i.e. the hot and runaway population defined in section 4.2) do not experience any radial transport (a) and where they experience a Rechester-Rosenbluth diffusion coefficient similarly to the thermal electrons (b). This allows us to distinguish the effect of radial transport of the superthermal electrons due to the magnetic perturbation. We also note that the radial transport of energetic electrons in the MeV-range might be lower than that given by the Rechester-Rosenbluth coefficient [68], due to the effect of particle orbit averaging over the fluctuation. Disabling the radial transport of the superthermal electrons may be seen as an upper estimate of the role of this phenomenon.

Without transport of the superthermal electrons, the hot-tail seed integrates to 1.15 MA. This value is around six orders of magnitude larger than the ~ 1 A seed expected from the tritium decay and Compton scattering during nuclear operation. Despite the logarithmic dependence of the final runaway current on the seed found in ref. [19], the hot-tail seed found here could therefore have a significant impact on the final runaway current. Moreover, the hot-tail seed found in ref. [19] to be sufficient to result in a final runaway current larger than 1 MA. Finally, the difference is even larger compared to the seed calculated for the two-stage injection studied earlier in figure 5.3. In this case, the hot-tail seed is practically suppressed, and is calculated to be at most $1.1 \cdot 10^{-12}$ A (still without transport of superthermal electrons)¹.

Looking at the radial profile in figure 5.6 a), we see that the hot-tail seed is strongly peaked at the center in this case. This is explained by the higher initial temperature and current density in the core compared to the rest of the plasma. The higher initial temperature increases the hot-tail generation by extending the tail of the initial momentum distribution. The effect of the higher current density is that it increases the local induced electric field during the disruption. We also see in figure 5.4 d) that the plateau-like phase of the temperature evolution directly after the injection persists for a shorter time in the core, which further favours hot-tail generation there, compared to the rest of the plasma.

When the transport of superthermals is included, we see in figure 5.6 b) that the runaway electrons are smeared out over a much wider part of the plasma. The total hot-tail seed is also significantly reduced. As the runaway electrons are transported out of the plasma as long as the magnetic field remains stochastic, the effective surviving magnitude of the hot-tail seed is dependent on the time it takes for the flux surfaces to re-heal. The maximum value of the total hot-tail seed is 43.6 kA, but after the maximum is passed it decays exponentially on a ~ 0.1 ms time scale. However, the inverse avalanche growth rate, which determines the avalanche *e*-folding time scale, is also of the order of 0.1-1 ms during a disruption in an ITER-like plasma [20]. Thus, the seed electrons might have time to contribute to a few *e*-folds before they are lost.

¹Such a small seed might be significantly affected by the numerical "leakage" of particles into the runaway population, which may occur when the upper momentum boundary $p_{\rm RE}$ is lower than the critical momentum for runaway acceleration. The calculated hot-tail seed should therefore be considered an upper estimate.



Figure 5.6: Hot-tail runaway seed current evolution for the single-stage injection studied in figure 5.4, with pellets consisting of 5% neon and 95% deuterium, with $2 \cdot 10^{24}$ particles, shattered into 1742 shards. The panels show results without transport of superthermals (a) and with transport of superthermals (b).

We now make a further study of the hot-tail seed generation for two-stage injections when varying the number of injected particles, using the same cases as in figure 5.5 (i.e. with $\langle v_{\rm p,D} \rangle = 800 \text{ m/s}$). As the purpose of this study is to compare the hot-tail generation between a single and a two-stage injection, we consider the simpler case without transport of superthermal electrons. Such a conservative case can be regarded as an upper limit of the hot-tail seed, the consideration of which is further motivated by the lack of detailed knowledge about the transport of superthermals. The resulting total hot-tail seeds are shown in figure 5.7 a). We note immediately that when core penetration with the deuterium injection is achieved, the seeds are several orders of magnitude smaller than those found for the single-stage injection discussed above. As expected, increasing the neon content strongly increases the hot-tail generation due to the resulting decrease in the thermal quench time. On the other hand, the general trend observed is that increasing the deuterium content, and hence the dilution following the first injection, can effectively reduce the hot-tail seed by several orders of magnitude.

Important to note, however, is that whether or not core penetration with the deuterium injection is achieved can alter the hot-tail seed by several orders of magnitude. When core penetration is not achieved, the core is rapidly cooled by neon radiation and transport, all the way from the initial temperature down to a few eV, without an intermediate tail equilibration. The situation in the core is therefore similar to the single-stage injection investigated in figure 5.6.

The radial profiles of the hot-tail seed for four values of $N_{\rm inj,D}$, each with $N_{\rm inj,Ne} = 10^{24}$, are shown in figure 5.7 b). When core penetration with the deuterium injection is not achieved (solid black curve), the profile peaks in the core, similarly to the single-stage case shown in figure 5.6 a). In contrast, the cases with larger deuterium pellets have their maxima close to the edge. As seen in figure 5.1, the temperature profiles following the deuterium injection for such cases are rather flat compared to the initial temperature profile, which lowers the relative efficiency of the hot-

tail generation in the core. Furthermore, the lower density, and hence the lower collisionality, makes the hot-tail generation at the edge relatively more efficient.



Figure 5.7: Maximum total hot-tail runaway seed as a function of the number of injected deuterium and neon particles (a) and the corresponding radial profiles of the hot-tail seed for $N_{\rm inj,Ne} = 10^{24}$ and varying $N_{\rm inj,D}$ (b). The parameters for the injections were $N_{\rm s,Ne} = 50$ for the neon shards, and the number of deuterium shards was chosen along the 97% assimilation contour in figure 5.1, when possible. This region is marked by the dotted grey line. To the left of the grey line, the deuterium pellets were shattered into 10 shards. A diffusive heat transport with $\delta B/B = 0.001$, constant in space and time, is activated once the neon shards enter the plasma.

5.3 Current quench simulations

We now turn our attention to the later stages of the disruptions and study the decay of the Ohmic current, as well as the remaining runaway generation and dissipation. These stages correspond to the green and blue shaded areas in figure 3.2. Scans are made over the same ranges of injection parameters used above in section 5.2.3 and 5.2.2 to study the hot-tail generation and transported thermal energy losses. At this stage, the thermal energy is almost completely dissipated, and the now low plasma temperature evolves slowly as the Ohmic heating varies during the current decay. When the rapid plasma cooling is complete, the stochastic flux surfaces re-heal, which we account for by switching of the transport due to magnetic perturbations at t = 16.4 ms. The hot-tail mechanism is also no longer active, and therefore our analysis can be done without resolving the kinetic distribution of the electrons. Instead, we now include the fluid-like sources described in section 4.4.2, adding further runaways to those resulting from the hot-tail seed calculated in section 5.2.3.

We consider separately the case where the tritium decay and Compton scattering seed mechanisms are active, corresponding to the nuclear operation phase, and the case without nuclear activity where these mechanisms are absent. The tritium decay and Compton scattering mechanisms generate a seed of the order of 0.1-1 A rather independently of the injection parameters, as discussed in section 3.3. Although seemingly small, these seeds may be multiplied by the avalanche mechanism to a final runaway current of several MA, as we will see shortly. During the current quench, when the plasma is relatively cold, the radiation transport properties of the plasma might have an important impact on the runaway generation. In particular, a preliminary estimate of opacity to Lyman radiation was presented in ref. [19], indicating that this effect could lower the runaway currents by up to several MA for large injected deuterium densities. Not only does the opacity impact the radiative losses, but it also affects the ionization and recombination rates (\mathcal{I}_{ij} and \mathcal{R}_{ij} in equation (4.8)) [69]. While the plasma is essentially transparent at most wavelengths, the opacity is significantly increased at the wavelength corresponding to resonant transitions [70]. This applies particularly to those transitions involving the ground state; as this is the state occupied by most ions and atoms, there are more atoms that can reabsorb a large fraction of the radiation at this wavelength.

The estimate shown in appendix B indicates that the plasma might only be transparent to a few percent of the Lyman radiation where the excited states are populated by excitations starting from the ground state, which primarily populates the lower excited states. This would have a significant impact on the contribution from the hydrogen isotopes to the radiation rate L_{ij} defined in section 4.2. On the other hand, the plasma is estimated to be transparent to the majority of the ionization/recombination radiation. A substantial part of this radiation consists of a continuum spectrum resulting from free-bound transitions, together with higher order Lyman lines resulting from the de-excitation of the high excited states thus populated. For such radiation, the plasma is largely transparent.

Moreover, it was indicated in ref. [71] that opacity to neon radiation is not expected to have a strong impact on disruption dynamics. We therefore consider the limiting cases where the plasma is assumed to be completely transparent or completely opaque to Lyman radiation, whilst being completely transparent to radiation from species other than hydrogen. For the completely transparent case, the radiation and ionization/recombination rates are taken from the ADAS database for all species as described in section 4.2. When the plasma is assumed to be opaque to Lyman radiation, the radiation and ionization/recombination rates for hydrogen species are instead taken from the AMJUEL database².

Figures 5.8 a-b) show the calculated runaway currents just preceding the runaway dissipation phase (corresponding to the boundary between the green and blue shaded areas in figure 3.2) for the different scenarios discussed above. This runaway current indicates the maximum runaway current that might strike the wall, depending on how far into the dissipation phase the plasma control is retained. The thin lines show the non-nuclear case, which is free from tritium decay and Compton scattering, while the thick lines show results for the nuclear case, i.e. with the tritium decay and Compton scattering seed mechanisms included. Panels c-d) show the corresponding Ohmic current quench times, defined as

$$t_{\rm CQ} = \frac{t(I_{\rm Ohm} = 0.2I_{\rm p}^{(t=0)}) - t(I_{\rm Ohm} = 0.8I_{\rm p}^{(t=0)})}{0.6}.$$
(5.1)

In panel a) and c), the plasma is assumed to be completely transparent, and in b) and d) the plasma is assumed to be completely opaque to Lyman radiation.

²http://www.eirene.de



Figure 5.8: Maximum runaway current (a-b) and current quench time (c-d) as functions of the number of injected deuterium and neon particles. The thick lines show the result for nuclear cases, with runaway generation by tritium decay and Compton scattering present. The thin lines show the result for non-nuclear cases, which are free from tritium decay and Compton scattering. In panel a) and c), the plasma was assumed to be completely transparent, while the plasma was assumed to be opaque to the Lyman lines in panels b) and d). The parameters for the injections were $N_{\rm s,Ne} = 50$ for the neon shards, and the number of deuterium shards was chosen along the 97% assimilation contour in figure 5.1, when possible. This region is marked by the dotted grey line. To the left of the grey line, the deuterium pellets were shattered into 10 shards. The magnitude of the magnetic perturbation was assumed to be $\delta B/B = 0.001$, activated once the neon shards enter the plasma, and switched off at t = 16.4 ms, when the rapid plasma cooling is complete.

A general trend seen in figure 5.8 a-b) is that the runaway currents are small or even negligible for the lowest considered $N_{inj,D}$ and $N_{inj,Ne}$, which do not give core penetration with the deuterium injection. However, for these cases the current quench times (above the axis in panels c-d)) are very long. The exact current quench times were not determined, as the simulations were stopped after 150 ms, but the decay rates indicate a time scale of the order of seconds. As $N_{inj,D}$ increases, the black curves in panels a-b), which initially indicated low runaway currents, first increases rapidly, before the increase either slows down significantly or turns to a decrease. Simultaneously, the current quench times in panels c-d) make a jump towards shorter times, and then decrease only slowly with increasing deuterium content. The reason for the initial jump is that, for low deuterium densities, the radiative cooling is not strong enough to overcome the ohmic heating. This leads to a reheating of the plasma from the 10 eV range to a few hundred eV, as the transport losses are no longer active. As the conductivity scales as $T_{\rm M}^{3/2}$ (see equation (4.9)), the re-heating greatly increases the conductivity, leading to a major reduction of the current decay rate. The increase in conductivity also reduces the induced electric field, and hence reduces the runaway generation. It should however be noted that the current quench times for these cases are much longer than the upper limit of 150 ms discussed in section 3.2.

For the rest of the parameter space, the current quench times are in the vicinity of the lower acceptable limit of 35 ms. Note, however, that in cases with a large runaway conversion, the runaway current aborts the current quench rather abruptly. The Ohmic current quench time calculated here is therefore a lower estimate of the current quench time relevant for the assessment of the electromagnetic loads due to eddy currents.

Another general trend seen in figure 5.8 is that at intermediate deuterium densities and high neon densities, the runaway currents decrease with increasing $N_{\rm inj,D}$. At the post-thermal quench temperatures of a few eV, the neon atoms are not fully ionized, while the deuterium ions remain practically fully ionized until the temperature drops below about 2 eV. Adding more deuterium therefore decreases the fraction of bound electrons. The presence of bound electrons results in a net enhancement of the avalanche, as discussed earlier in section 3.3, and therefore a reduced fraction of bound electrons reduces the runaway generation.

As the Ohmic current decays the Ohmic heating decreases, and eventually the temperature falls below 2 eV so that the deuterium starts to recombine. At high deuterium densities, the increased radiative losses can cause this to happen while there is still a substantial part of the Ohmic current left that can be converted to a runaway current. When the deuterium recombines, the fraction of bound electrons increases and enhances the avalanche. This is the reason behind the increase in the runaway current at high deuterium densities in figure 5.8 a). How much current remains when the deuterium starts to recombine is however dependent on the radiation transport properties of the plasma. When the deuterium starts to recombine, it also starts to contribute significantly to the radiation, and therefore opacity to Lyman radiation considerably reduces the radiative losses. This reduction in the radiative losses means that the Ohmic current must be smaller for the deuterium to recombine, explaining the lower maximum runaway currents at high deuterium densities in figure 5.8 b). The effect of the radiative properties of the plasma are discussed in more detail later in this section.

Finally, we see a clear difference when runaway seed generation from tritium decay and Compton scattering is included, compared to when it is not. The differing factor is however far from the several orders of magnitude difference between the 0.1-1 A seed produced by the tritium decay and Compton scattering and the hot-tail seeds shown in figure 5.7. This is due to the logarithmic dependence of the maximum runaway current on the seed found in ref. [19], which was also discussed in section 5.2.3. This weak dependence is explained by the self-regulating interaction between the runaway current and the electric field; when the runaway current becomes comparable to the

case	$N_{\rm inj,D}$	$N_{\rm inj,Ne}$
1	$0.5 \cdot 10^{24}$	10^{25}
2	$0.3\cdot10^{24}$	10^{22}
3	$2\cdot 10^{24}$	10^{24}
4	$1\cdot 10^{24}$	10^{22}

Table 5.1: Representative cases used to illustrate the four qualitatively different regions found in the $N_{inj,D} - N_{inj,Ne}$ -space.

remaining Ohmic current, the induced electric field is reduced, which reduces the avalanche growth rate. Notably, however, despite this weak dependence, we see a significant difference between the nuclear and non-nuclear cases. In the non-nuclear cases, we found maximum runaway currents lower than the critical 2 MA limit discussed in section 3.2, which also had acceptable current quench times. In the nuclear cases, on the other hand, no such cases were found. The implications of these findings are discussed later in chapter 6.

5.3.1 Current quench and runaway dynamics

For the remainder of the analysis in this section, we identify four qualitatively different regions of the parameter space: (1) a region with high injected neon density, with strong runaway generation, where the runaway current decreases with increasing deuterium density; (2) a region at low injected neon and deuterium densities, where core penetration is not achieved with the deuterium injection, giving low runaway generation, but current quench times exceeding the limit; (3) a region at high injected deuterium densities with strong runaway generation, that increases with increasing deuterium density; (4) a region intermediate between (1) and (3) where the runaway current has its minimum in the part of parameter space with acceptable current quench times. These four regions are similar to the ones found with the fluid model GO used in ref. [19], when scanning over prescribed neon and deuterium densities with a flat profile. In this section, we illustrate these regions by detailing one representative case from each, given in table 5.1. This is done for the simulations assuming a completely transparent plasma, because the features of interest are more clearly pronounced with this assumption. It should however be noted that the estimates in appendix B indicate that the more optimistic results obtained assuming complete opacity to Lyman radiation might be closer to reality.

The time evolution of the most relevant plasma parameters for case 1 are shown in figure 5.9. Panels a), b) and c) show the temperature, electric field, and runaway current density during the final phase of the runaway generation. The temperature stabilises at a few eV, with a rather flat profile in most of the plasma, with a peak in the centre. The induced electric field is of the order of a few tens of V/m, and the current quench time is a few tens of milliseconds. After about 25 ms, the Ohmic current is completely dissipated and replaced by a runaway current of 5.2 MA, as shown in panel d). The profile of the runaway current has its peak about 80 cm outside the core, coinciding with the peak of the deuterium and neon density profiles. The subsequent dissipation of the runaway current is rather slow for the relatively



modest electron density prevailing in this case.

Figure 5.9: Spatio-temporal evolution of a) temperature, b) electric field and c) runaway current density for case 1, defined in table 5.1, during the final part of the runaway generation phase. The times indicated are those elapsed since the start of the first deuterium injection. d) shows the time evolution of the total (black solid) and runaway (green dashed) current.

The time evolution of the temperature, electric field and runaway current for cases 2-4 are shown in figure 5.10. For case 2, shown in panels a), d) and g), we clearly see a re-heating approaching 200 eV. In the regions where the plasma is re-heated, the electric field is relatively low, and the runaway generation is negligible in all parts of the plasma. The resulting evolution of the total Ohmic current, shown in figure 5.11 a), shows an initial relatively fast drop when the current in the outer part of the plasma (that is not re-heated) dissipates. After that, the remaining current dissipates at a very slow rate, with a current quench time of the order of seconds.



Figure 5.10: Time evolution of the temperature (a,b and c), electric field (d,e and f) and runaway current density (g,h and j) for the representative cases 2-4 of the $N_{\rm inj,D} - N_{\rm inj,Ne}$ parameter space, defined in table 5.1. The times indicated are those elapsed since the start of the first deuterium injection. Left column shows case 2, middle column case 3, and right column case 4.



Figure 5.11: a) Time evolution of the total plasma current (black) and runaway current (blue) for case 2 (solid), case 3 (dash-dotted) and case 4 (dashed) defined in table 5.1. b) initial current profile (black solid) and final runaway current profiles for case 3 (dash-dotted) and case 4 (dashed).

In case 3, shown in figure 5.10 b), e) and h), the temperature drops close to 1 eV over most of the plasma while there is still a substantial Ohmic current. The deuterium starts to recombine, and the local electric field initially increases and is then dissipated over a relatively short time scale of a few milliseconds. In the time between the recombination of the deuterium and the decay of the electric field, the local avalanche generation is very strong. This results in a large runaway current of 7.8 MA, with profile peaking in the outer region where the temperature first drops to 1 eV. The relatively high injected deuterium and neon densities result in a comparatively fast decay of the runaway current in case 3, as seen in figure 5.11 a), and the runaway current eventually striking the wall is therefore somewhat dependent on how long the plasma position can be controlled.

Finally, case 4, shown in figure 5.10 c), f) and i), can be seen as an intermediate case between case 1 and 3. The deuterium density is high enough to reduce the enhancement of the avalanche by the bound electrons in the neon ions, but not so high as to cool the plasma to the temperature where deuterium recombines, before the runaway generation phase is close to finished. The resulting runaway current is 3.9 MA, which is the lowest runaway current found for a completely transparent plasma in the region of parameter space with potentially acceptable current quench times. The decay rate of the runaway current however is quite modest compared to case 3. The injection parameters resulting in the smallest runaway current striking the wall is therefore somewhat dependent on the time during which the plasma position remains stable. It should be noted, however, that a faster current decay makes it harder to maintain control of the plasma position [26]. As a result, the plasma may strike the wall at an earlier stage, which reduces the impact of the deposition, and favours the case with the lowest maximum runaway current.

5.3.2 Qualitative interpretation based on energy balance

The differences in the dynamics between the four cases above are primarily caused by differences in the balance between Ohmic heating and radiative losses during the current quench. This balance can be qualitatively understood by comparing Ohmic heating to radiative losses at different temperatures, assuming an equilibrium distribution of charge states (which may be calculated using equation (4.7)). Such a comparison is shown in figure 5.12 for the above four cases, considering the volume averaged final deuterium and neon densities. The solid black line shows the radiative losses assuming a completely transparent plasma, calculated by the second term on the right hand side of equation (4.8). The green dashed line shows the corresponding radiative losses when the plasma is assumed to be opaque to Lyman radiation. The Ohmic heating is calculated using the first term on the right hand side of equation (4.8), for two different current densities: 5 MA/m^2 (blue dashed), taken as a representative value for the maximum current density at the radial peak, similar to that seen in figure 5.3 f), and 0.5 MA/m^2 (blue dotted), representing the phase where the Ohmic current has partially, but not completely, decayed.

In case 1, the only equilibrium temperature (i.e. where the heating and losses balance each other) is of the order of a few eV for both considered Ohmic current densities. The plasma will therefore equilibrate at a temperature in this range throughout the current quench. At these temperatures, the hydrogen isotopes remain fully ionized. However, the comparatively large number of bound electrons in the neon ions enhances the avalanche, making the maximum runaway current relatively large.

In case 2 on the other hand, there is another equilibrium at about 200 eV for $j_{\rm Ohm} = 5 \text{ MA/m}^2$, in addition to the one at about 10 eV. Depending on the plasma parameters, when the magnetic perturbation is turned off, parts of the plasma may therefore equilibrate at about 200 eV. As there is also an equilibrium in the few eV range, the regions with a lower Ohmic heating will equilibrate at such a temperature. In those regions, the local current density will decay at a similar rate as in case 1, resulting in the initial current drop in figure 5.11. However, in the parts evolving towards the ~ 200 eV equilibrium, the conductivity remains comparatively large, leading to a slow current decay, low electric fields and weak runaway generation.

In case 3, the large deuterium density increases the radiative losses at lower temperatures, especially below ~ 2 eV, where a substantial fraction of the deuterium recombines. When this happens, the deuterium starts to contribute directly to the radiation rather than just increasing the electron density. For $j_{\rm Ohm} = 0.5 \text{ MA/m}^2$, and a transparent plasma, the equilibrium temperature is now close to 1 eV, where the ionized fraction of the deuterium is only a few percent. This results in an enhancement of the avalanche, and the low temperature also favours larger induced electric fields, albeit with a faster local decay rate. Moreover, note that although the Ohmic current density is somewhat modest in the ~ 1 eV regions, the local conversion from Ohmic to runaway current is not limited by the local Ohmic current density. The current from other parts of the plasma can diffuse into the cold regions, potentially causing a local increase in the current density. Altogether, these effects give rise to the increase in the maximum runaway current at large deuterium densities seen in figure 5.8.

If the plasma is assumed to be opaque to Lyman radiation, however, the radiative losses at low temperatures are considerably reduced. This is particularly true for the line radiation following excitations from the ground state. The remaining increase in the radiative losses below $\sim 2 \text{ eV}$ is instead primarily due to ionization/recombination radiation of hydrogen species. The drop to $\sim 1 \text{ eV}$ is now postponed until the Ohmic current has decreased further than in the transparent case. Although the Lyman opaque case is not qualitatively different to the transparent one, the postponed temperature drop to $\sim 1 \text{ eV}$ reduces the maximum runaway current at large deuterium densities, and shifts the increasing trend to higher deuterium densities.

Finally, case 4 can be regarded as a compromise between the features characteristic of the other cases. On the one hand, the amount of deuterium and neon is large enough to avoid the re-heating seen in case 2 and to avoid the enhancement of the avalanche by electrons bound to neon ions seen in case 1. On the other hand, it is not large enough to cause a final temperature drop to ~ 1 eV before the remaining Ohmic current has fallen. As a result, the current quench time is comparable to the lower limit of 35 ms, and the maximum runaway current is relatively modest.



Figure 5.12: Radiative power loss as a function of temperature, for a transparent plasma (solid black) and for a plasma opaque to Lyman radiation (dash-dotted green), compared to the Ohmic heating calculated for $j_{\rm Ohm} = 5 \text{ MA/m}^2$ (dashed blue) and $j_{\rm Ohm} = 0.5 \text{ MA/m}^2$ (dotted blue). The panels show case 1 (a), case 2 (b), case 3 (c) and case 4 (d), defined in table 5.1, using the volume averaged final deuterium and neon densities and an equilibrium distribution of charge states.

6

Concluding remarks

Disruptions are a severe threat to the future of fusion energy based on the tokamak design. A reliable disruption mitigation system is of utmost importance for the success of future large, high-current devices such as ITER. To this end, ITER plans for a disruption mitigation system based on Shattered Pellet Injections (SPI). In this work, we have undertaken numerical simulations to assess the performance of the most up to date SPI schemes, covering a wide range of injection parameters. The numerical simulations were performed with the DREAM (Disruption Runaway Electron Avoidance Model) code, which was extended within this work with an SPI model based on the Neutral Gas Shielding (NGS) model for pellet ablation.

Our assessment of the performance of the disruption mitigation system concerns three of the main issues related to disruptions: mitigation of the localised heat loads and electromagnetic forces on the vessel, and the reduction of the runaway current generation. The mitigation of the localised heat loads is achieved by isotropically dissipating the thermal energy content through radiation. The performance in this regard was quantified by the fraction of the initial thermal energy content conducted to the wall rather than radiated. The aim is a transported fraction lower than 10%. To avoid excessive electromagnetic forces, the current quench time must be long enough to sufficiently reduce the eddy currents induced in the structures surrounding the plasma, but short enough to avoid excessive halo currents. For ITER, the allowed range for the current quench time is therefore between 35 and 150 ms. The current quench time scale is essentially determined by the post-thermal quench temperature, which is, in turn, determined by the radiation losses and Ohmic heating. Finally, the runaway current impacting the wall should be lower than 2 MA. Given any initial runaway seed current, the main source of runaway generation is the avalanche mechanism. This is sensitive to the electric field and distribution of ionization stages during the current quench, with a high electric field and a large fraction of bound electrons giving a high avalanche generation.

6.1 Summary of results

Our results effectively address the three main requirements of the disruption mitigation system mentioned above. This was done through simulations of a two-stage SPI scheme, starting with dilution cooling by a deuterium injection followed by a radiatively cooling neon injection. We found that such an injection scheme significantly impedes the thermal energy transport due to magnetic perturbations, which can reduce localised heat loads. This is explained by the decrease in the thermal motion due to the dilution cooling in the first stage, before the stochastisation of the magnetic field, which reduces the cross-field transport along the stochastic magnetic field lines.

If the injected neon and deuterium densities are too low, the current quench time becomes longer than the upper limit for avoiding unacceptable halo currents. The reason for this was the resulting incomplete cooling of the plasma, and the associated high conductivity in the parts of the plasma remaining at a comparatively high temperature. For large injected densities, however, the simulated current quench times were found to be close to the lower acceptable limit, with a rather weak dependence on the injected densities.

Another major concern related to the rapid cooling in disruptions has been the potential generation of a strong hot-tail seed. We found that the above mentioned two-stage SPI scheme can effectively reduce the hot-tail runaway seed generation by several orders of magnitude, enabling a final runaway current below the acceptable limit of 2 MA in the absence of other significant seed mechanisms. This reduction can be explained by the intermediate equilibration of the hot tail of the electron momentum distribution between the injections. Notably, this strong reduction was obtained despite the conservative assumption that the superthermal electrons were not affected by the magnetic field perturbations.

For the final runaway current, a non-monotonic dependence on the injected deuterium density was found. At low deuterium densities, an increased deuterium density was found to reduce the avalanche, while at large deuterium densities, the avalanche was found to be enhanced by the higher electric fields and substantial recombination, resulting from the strong radiative cooling. The increase in the runaway current for large deuterium densities favours an intermediate deuterium injection of about $1 - 2 \cdot 10^{24}$ atoms, depending on the opacity of the plasma to Lyman radiation. This sets an upper limit on the dilution cooling that may be employed. Our results indicate, however, that the deuterium injections giving the minimum runaway current, followed by a neon injection in the range of $10^{22} - 10^{24}$ atoms, gives current quench times and transported fractions of the thermal energy within, or at least in the vicinity of, the acceptable ranges.

For the non-nuclear cases, these parameters were also found to give peak runaway currents lower than the acceptable limit of 2 MA, especially when assuming the plasma to be opaque to Lyman radiation. For the nuclear cases, on the other hand, no such scenarios were found, which is quite alarming for the nuclear phase of the ITER operation. It should however be noted that the lowest calculated peak runaway currents of around 4 MA, again somewhat dependent on the opacity properties of the plasma, are of the same order of magnitude as the acceptable limit of 2 MA. The runaway current remaining upon wall impact might also be lower than the peak runaway current evaluated here due to dissipation, depending on how long the plasma position remains stable. The dissipation might also be significantly affected by the presence of magnetic perturbations remaining during the current quench and runaway dissipation phase, which was not considered in this work. The remaining necessary reduction of the runaway current might be further accomplished by involving additional mitigation techniques, such as runaway dissipation by deliberately applied plasma perturbations during the current quench and runaway dissipation phase.

6.2 Outlook

Although the simulations performed in this work are based on ITER-like parameters, the qualitative features of the results might be of interest for the design of disruption mitigation systems at other devices such as e.g. SPARC. One of the main advantages of a two-stage SPI scheme found in this work is the efficient suppression of the hot-tail generation. Previous studies modeling runaway generation in the SPARC tokamak indicated that the presence of a hot-tail seed can lead to several MAs of runaway currents [9], while the runaway generation in the absence of a hot-tail seed, but otherwise similar conditions, is of the order of 0.1 MA [35]. For this reason, a disruption mitigation scheme which effectively suppresses the hot-tail generation might be of interest for SPARC, motivating similar studies to those performed in this work with SPARC-like parameters.

The benefits of a two-stage SPI scheme indicated by our results also motivate more quantitatively accurate studies, with a more rigorous treatment of the most simplified aspects of the present model. While the model used in this work comprises an integrated framework accounting for many of the relevant aspects of disruption mitigation by SPI, the various aspects are treated with different levels of sophistication. The evolution of the current and momentum distribution of the electrons are treated rigorously in the present model. On the other hand, some aspects of the transport properties of the thermal energy and plasma ions, as well as the evolution of the pellet material, are treated in a simplified manner.

The SPI model used in this work contains a number of simplifications. One such simplification is the use of the NGS ablation model, which is based on a simplified spherical pellet shard geometry. In addition, it neglects the details of the ablating electron momentum distribution, as well as the electrostatic and diamagnetic shielding described in section 3.5.1. Nevertheless, the NGS model has been shown to compare reasonably well with experiments, and the effect of the above simplifications have been estimated in the literature to counteract each other [43]. Therefore, while a more advanced ablation model could result in an order unity correction to the deposition profile, the discrepancy compared to the NGS model is expected to be quite modest.

A larger correction might be caused by the details of the expansion process of the pellet material between the ablation and final deposition. In the present model, this process has simply been assumed to be instantaneous and local at the flux surfaces. In reality, however, the homogenisation and equilibration process takes of the order of 1 ms, as discussed in section 3.5.2. This time scale is similar to the time it takes the plume of shards to pass a given flux surface, and can therefore be important to account for. The process also covers several flux surfaces, as the newly ablated material is subject to an $E \times B$ -drift towards the low field side.

Apart from the direct impact on the timing and position of the density increase, the features of the expansion process might significantly alter the ablation. For pellets with a significant impact on the plasma temperature, the interaction between the plasma and the pellet material is self-regulating; when the pellet material is ablated, the plasma is cooled, at first primarily by dilution, resulting in a slowing down of the ablation. The finite expansion time and drifts reduce this self-regulation by delaying the plasma response, and shifting it away from the position of the shards. The ablation would thus be faster, so that more material is deposited earlier along the shard trajectories.

The impact of the self-regulation is estimated in appendix C, by introducing an artificial shift in the deposition kernel $H(r, \rho_{p,k})$ in equation (4.3). These estimates indicate that the self-regulation substantially affects the final density profile. However, a finite expansion time means that the thermal energy absorbed in the ablation cloud surrounding the shards is not immediately returned to the background plasma. Depending on the size of the ablation cloud, this heat absorption might compensate for the reduced dilution cooling at the position of the pellet shards. In such a case, it is estimated in appendix C that the final density profile might be rather similar to that obtained with the model used in this work, which neglects the finite expansion time and drifts. A quantitatively accurate model would however have to account for the features of the expansion process self-consistently.

A second area of simplification employed in the present work concerns the geometry and interaction with the structures surrounding the plasma. The geometrical simplifications include the circular plasma cross section, neglect of the toroidicity and the assumption of flux surface homogenised quantities. Relaxing these assumptions would allow for modelling of transient 3D features of the plasma profiles, as well as introducing geometrical order unity corrections to the transport processes involved. For instance, elongating the plasma increases the cross sectional area. This leads to an increase of the time scale for diffusive transport across the plasma cross section, and also decreases the density for a given number of deposited particles. Regarding the surrounding structures, their geometry and conductive properties introduce corrections to the electric field boundary condition. Support for a shaped geometry (although still with constant plasma parameters over the flux surfaces) and a finite wall conductivity has recently been implemented in DREAM, and the sensitivity to these features could therefore be studied in the future.

Finally, another simplification in the present model is the prescribed evolution of the magnetic perturbations. The prescribed magnetic perturbation is sufficient to study qualitative trends involving transport due to magnetic perturbations, as done in this work. However, self-consistent and quantitatively accurate simulations would require coupling to a magnetohydrodynamic (MHD) model, such as JOREK [24], for the evolution of the magnetic field. The evolution of the magnetic field through the current quench might also be of interest, potentially including deliberately induced perturbations to increase the runaway dissipation, as discussed above. Even the comparatively smaller magnetic perturbations that may be present during the current quench can have a relatively large impact on the runaway generation and dissipation [72]. This applies especially for cases with an off-center runaway current profile, as magnetic perturbations induced during the current quench are expected to only partially penetrate into the plasma. However, it may apply also for core-centered runaway current profiles given that the perturbations can cover a major part of the plasma where the runaway current resides.

A major issue with the use of MHD models such as JOREK is the significant

amount of computational resources required. A full scale 3D simulation of a thermal quench with JOREK run on a large scale computing cluster takes of the order of months to finish [73]. The thermal quench simulations shown in this work, on the other hand, take at most a couple of days on a desktop, when invoking the kinetic equation including transport of superthermal electrons. The run-times for pure fluid simulations, ran through both the thermal and current quench, are no longer than a couple of minutes. The orders of magnitude lower computational expense substantially increases the feasibility of exploring a wide range of injection parameters. Instead of attempting a full scale coupling of the evolution of the fluid, kinetic and MHD properties of the plasma, a more feasible approach might be to use the output of a full scale MHD model such as JOREK to design a simplified MHD model. Such a model could then be efficiently integrated in a framework like the one used in this work. If successful, such a model could significantly improve the disruption modelling accuracy, as the coupling of the MHD, kinetic and fluid properties of the plasma might be one of the main questions remaining in order to make self-consistent, quantitatively accurate, predictions for the performance of disruption mitigation systems.

6. Concluding remarks

Bibliography

- H. Cabal, Y. Lechon, C. Bustreo, F. Gracceva, M. Biberacher, D. Ward, D. Dongiovanni, and P. E. Grohnheit, "Fusion power in a future low carbon global electricity system", *Energy Strategy Reviews*, vol. 15, 2017. DOI: 10.1016/j.esr.2016.11.002.
- [2] J. P. Freidberg, *Plasma physics and fusion energy*, 1st ed. Cambridge, UK: Cambridge university press, 2008.
- [3] F. Chen, An indispensable truth: how fusion power can save the planet. New York, NY, USA: Springer Science & Business Media, 2011.
- [4] F. Chen, Introduction to Plasma Physics and Controlled Fusion, 1st ed. New York, NY, USA: Springer International Publishing, 1974.
- [5] E. Hollmann, P. Aleynikov, T. Fülöp, D. Humphreys, V. Izzo, M. Lehnen, V. Lukash, G. Papp, G. Pautasso, F. Saint-Laurent, *et al.*, "Status of research toward the ITER disruption mitigation system", *Physics of Plasmas*, vol. 22, 021802–1, 2015. DOI: 10.1063/1.4901251.
- [6] A. J. Creely, M. J. Greenwald, S. B. Ballinger, D. Brunner, J. Canik, J. Doody, T. Fülöp, D. T. Garnier, R. Granetz, T. K. Gray, and et al., "Overview of the SPARC tokamak", *Journal of Plasma Physics*, vol. 86, p. 865860502, 2020. DOI: 10.1017/S0022377820001257.
- [7] UK Atomic Energy Authority, https://step.ukaea.uk/.
- [8] ITER organization, http://www.iter.org.
- [9] R. Sweeney, A. J. Creely, J. Doody, T. Fülöp, D. T. Garnier, R. Granetz, M. Greenwald, L. Hesslow, J. Irby, V. A. Izzo, and et al., "MHD stability and disruptions in the SPARC tokamak", *Journal of Plasma Physics*, vol. 86, p. 865 860 507, 2020. DOI: 10.1017/S0022377820001129.
- [10] P. Helander and D. Sigmar, Collisional Transport in Magnetized Plasmas. Cambridge, UK: Cambridge University Press, 2005.
- O. Embréus, S. Newton, A. Stahl, E. Hirvijoki, and T. Fülöp, "Numerical calculation of ion runaway distributions", *Physics of Plasmas*, vol. 22, 5 2015.
 DOI: 10.1063/1.4921661.
- H. Dreicer, "Electron and ion runaway in a fully ionized gas I", *Physical Review*, vol. 115, p. 238, 1959. DOI: 10.1103/PhysRev.115.238.

- P. Helander, H. Smith, T. Fülöp, and L.-G. Eriksson, "Electron kinetics in a cooling plasma", *Physics of Plasmas*, vol. 11, p. 5704, 2004. DOI: 10.1063/1. 1812759.
- H. Smith, P. Helander, L.-G. Eriksson, and T. Fülöp, "Runaway electron generation in a cooling plasma", *Physics of Plasmas*, vol. 12, p. 122505, 2005.
 DOI: 10.1063/1.2148966.
- [15] H. M. Smith and E. Verwichte, "Hot tail runaway electron generation in tokamak disruptions", *Physics of Plasmas*, vol. 15, p. 072502, 2008. DOI: 10.1063/1.2949692.
- [16] J. R. Martín-Solís, A. Loarte, and M. Lehnen, "Formation and termination of runaway beams in ITER disruptions", *Nuclear Fusion*, vol. 57, p. 066 025, 2017. DOI: 10.1088/1741-4326/aa6939.
- [17] Y. A. Sokolov, ""multiplication" of accelerated electrons in a tokamak", JETP Lett. (USSR) (Engl. Transl.); (United States), vol. 29, 1979.
- [18] M. Rosenbluth and S. Putvinski, "Theory for avalanche of runaway electrons in tokamaks", *Nuclear Fusion*, vol. 37, p. 1355, 1997. DOI: 10.1088/0029-5515/37/10/i03.
- [19] 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, p. 475 860 401, 2020. DOI: 10.1017/S0022377820000859.
- [20] L. Hesslow, O. Embréus, O. Vallhagen, and T. Fülöp, "Influence of massive material injection on avalanche runaway generation during tokamak disruptions", *Nuclear Fusion*, vol. 59, p. 084 004, 2019. DOI: 10.1088/1741-4326/ab26c2.
- [21] M. Lehnen, S. Jachmich, U. Kruezi, and the ITER DMS task force, *The ITER disruption mitigation strategy*, https://conferences.iaea.org/event/217/contributions/17867/, Presented by M. Lehnen at the IAEA Technical Meeting on Plasma Disruption and their Mitigation, 2020.
- [22] S. Chiu, M. Rosenbluth, R. Harvey, and V. Chan, "Fokker-planck simulations mylb of knock-on electron runaway avalanche and bursts in tokamaks", *Nuclear Fusion*, vol. 38, p. 1711, 1998. DOI: 10.1088/0029-5515/38/11/309.
- [23] R. W. Harvey, V. S. Chan, S. C. Chiu, T. E. Evans, M. N. Rosenbluth, and D. G. Whyte, "Runaway electron production in DIII-D killer pellet experiments, calculated with the CQL3D/KPRAD model", *Physics of Plasmas*, vol. 7, p. 4590, 2000. DOI: 10.1063/1.1312816.
- [24] E. Nardon, D. Hu, M. Hoelzl, and D. B. and, "Fast plasma dilution in ITER with pure deuterium shattered pellet injection", *Nuclear Fusion*, vol. 60, p. 126040, 2020. DOI: 10.1088/1741-4326/abb749.
- [25] M. Hoppe, O. Embreus, and T. Fülöp, DREAM: a fluid-kinetic framework for tokamak disruption runaway electron simulations, 2021. arXiv: 2103.16457 [physics.plasm-ph].

- T. Hender, J. Wesley, J. Bialek, A. Bondeson, A. Boozer, R. Buttery, A. Garofalo, T. Goodman, R. Granetz, Y. Gribov, O. Gruber, M. Gryaznevich, G. Giruzzi, S. Günter, N. Hayashi, P. Helander, C. Hegna, D. Howell, D. Humphreys, G. Huysmans, A. Hyatt, A. Isayama, S. Jardin, Y. Kawano, A. Kellman, C. Kessel, H. Koslowski, R. L. Haye, E. Lazzaro, Y. Liu, V. Lukash, J. Manickam, S. Medvedev, V. Mertens, S. Mirnov, Y. Nakamura, G. Navratil, M. Okabayashi, T. Ozeki, R. Paccagnella, G. Pautasso, F. Porcelli, V. Pustovitov, V. Riccardo, M. Sato, O. Sauter, M. Schaffer, M. Shimada, P. Sonato, E. Strait, M. Sugihara, M. Takechi, A. Turnbull, E. Westerhof, D. Whyte, R. Yoshino, H. Zohm, D. the ITPA MHD, and M. C. T. Group, "Chapter 3: MHD stability, operational limits and disruptions", *Nuclear Fusion*, vol. 47, p. 128, 2007. DOI: 10.1088/0029-5515/47/6/S03.
- [27] 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, p. 475 810 502, 05 2015. DOI: 10.1017/S0022377815000513/.
- [28] O. Embréus, A. Stahl, and T. Fülöp, "Effect of bremsstrahlung radiation emission on fast electrons in plasmas", New Journal of Physics, vol. 18, p. 093 023, 2016.
- [29] J. Wesson, *Tokamaks*, 4th ed. Oxford, UK: Oxford university press, 2011.
- [30] P. Svensson, "Effects of magnetic perturbations and radiation on the runaway avalanche", Master's thesis, 2020.
- [31] D. Ward and J. Wesson, "Impurity influx model of fast tokamak disruptions", Nuclear Fusion, vol. 32, p. 1117, 1992. DOI: 10.1088/0029-5515/32/7/i03.
- [32] L. Hesslow, L. Unnerfelt, O. Vallhagen, O. Embréus, M. Hoppe, G. Papp, and T. Fülöp, "Evaluation of the Dreicer runaway growth rate in the presence of high-Z impurities using a neural network", *Journal of Plasma Physics*, vol. 85, p. 475 850 601, 2019. DOI: 10.1017/S0022377819000874.
- [33] O. Linder, E. Fable, F. Jenko, G. Papp, and G. Pautasso, "Self-consistent modeling of runaway electron generation in massive gas injection scenarios in ASDEX upgrade", *Nuclear Fusion*, vol. 60, p. 096 031, 2020. DOI: 10.1088/ 1741-4326/ab9dcf.
- [34] K. I. Björk, O. Vallhagen, G. Papp, C. Reux, O. Embreus, E. Rachlew, T. Fülöp, the ASDEX Upgrade Team, J. contributors, and the EUROfusion MST1 Team, *Modelling of runaway electron dynamics during argon-induced disruptions in ASDEX Upgrade and JET*, 2021. arXiv: 2101.02575 [physics.plasm-ph].
- [35] T. Fülöp, P. Helander, O. Vallhagen, O. Embréus, L. Hesslow, P. Svensson, A. J. Creely, N. T. Howard, and P. Rodriguez-Fernandez, "Effect of plasma elongation on current dynamics during tokamak disruptions", *Journal of Plasma Physics*, vol. 86, p. 474 860 101, 2020. DOI: 10.1017/S002237782000001X.
- [36] P. Aleynikov and B. N. Breizman, "Generation of runaway electrons during the thermal quench in tokamaks", *Nuclear Fusion*, vol. 57, p. 046009, 2017. DOI: 10.1088/1741-4326/aa5895.

- [37] O. Embréus, A. Stahl, and T. Fülöp, "On the relativistic large-angle electron collision operator for runaway avalanches in plasmas", *Journal of Plasma Physics*, vol. 84, p. 905 840 102, 2018. DOI: 10.1017/S002237781700099X.
- [38] G. Papp, G. Pautasso, J. Decker, M. Gobbin, P. Mccarthy, P. Blanchard, D. Carnevale, D. Choi, S. Coda, B. Duval, R. Dux, B. Erdös, B. Esposito, O. Ficker, R. Fischer, C. Fuchs, C. Galperti, L. Giannone, A. Gude, B. Labit, K. Lackner, T. Lunt, L. Marelli, P. Martin, A. Mlynek, M. Maraschek, P. Marmillod, M. Nocente, Y. Peysson, P. Piovesan, V. Plyusnin, G. Pokol, P. Poloskei, S. Potzel, C. Reux, F. Saint-Laurent, O. Sauter, B. Sieglin, U. Sheikh, C. Sommariva, W. Suttrop, G. Tardini, D. Testa, W. Treutterer, M. Valisa, T. Team, and E. M. Team, "Runaway electron generation and mitigation on the european medium sized tokamaks ASDEX Upgrade and TCV", in *Proceedings* of the 2016 IAEA Fusion Energy Conference, 2016.
- [39] 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. M. D. Aguilera, G. Matthews, J. Mlynář, E. Nardon, E. Nilsson, C. P. von Thun, V. Riccardo, F. Saint-Laurent, A. Shevelev, G. Sips, and C. S. and, "Runaway electron beam generation and mitigation during disruptions at JET-ILW", *Nuclear Fusion*, vol. 55, p. 093013, 2015. DOI: 10.1088/0029-5515/55/9/093013.
- [40] B. Pégourié and J. Picchiottino, "Plasma density buildup after pellet injection", *Physics of Plasmas*, vol. 3, p. 4594, 1996. DOI: 10.1063/1.872030.
- [41] S. Milora, W. Houlberg, L. Lengyel, and V. Mertens, "Pellet fuelling", Nuclear Fusion, vol. 35, p. 657, 1995. DOI: 10.1088/0029-5515/35/6/i04.
- B. N. Breizman, P. Aleynikov, E. M. Hollmann, and M. Lehnen, "Physics of runaway electrons in tokamaks", *Nuclear Fusion*, vol. 59, p. 083001, 2019. DOI: 10.1088/1741-4326/ab1822.
- B. Pégourié, "Review: Pellet injection experiments and modelling", *Plasma Physics and Controlled Fusion*, vol. 49, p. 87, 2007. DOI: 10.1088/0741-3335/49/8/r01.
- [44] 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", *The Physics of Fluids*, vol. 21, p. 1735, 1978. DOI: 10.1063/1.862088.
- B. Pegourie, J.-M. Picchiottino, H.-W. Drawin, A. Geraud, and M. Chatelier, "Pellet ablation studies on Tore Supra", *Nuclear Fusion*, vol. 33, p. 591, 2002. DOI: 10.1088/0029-5515/33/4/106.
- [46] L. Lengyel, K. Büchl, G. Pautasso, L. Ledl, A. Ushakov, S. Kalvin, and G. Veres, "Modelling of impurity pellet ablation in ASDEX upgrade (neon) and Wendelstein W7-AS (carbon) by means of a radiative ('killer') pellet code", *Nuclear Fusion*, vol. 39, p. 791, 1999. DOI: 10.1088/0029-5515/39/6/307.

- [47] D. Morozov, V. Gervids, I. Senichenkov, I. Veselova, V. Rozhansky, and R. Schneider, "Ionization–recombination processes and ablation cloud structure for a carbon pellet", *Nuclear Fusion*, vol. 44, p. 252, 2004. DOI: 10.1088/0029-5515/44/2/005.
- [48] A. MacAulay, "Geometrical, kinetic and atomic physics effects in a two dimensional time dependent fluid simulation of ablating fuel pellets", *Nuclear Fusion*, vol. 34, p. 43, 1994. DOI: 10.1088/0029-5515/34/1/i03.
- [49] P. B. Parks and L. R. Baylor, "Effect of parallel flows and toroidicity on cross-field transport of pellet ablation matter in tokamak plasmas", *Phys. Rev. Lett.*, vol. 94, p. 125002, 2005. DOI: 10.1103/PhysRevLett.94.125002.
- [50] P. Parks, J. Leffler, and R. Fisher, "Analysis of low Z impurity pellet ablation for fusion diagnostic studies", *Nuclear Fusion*, vol. 28, p. 477, 1988. DOI: 10.1088/0029-5515/28/3/012.
- [51] P. Parks, A theoretical model for the penetration of a shattered-pellet debris plume, https://tsdw.pppl.gov/Talks/2017/Lexar/Wednesday%
 20Session%201/Parks.pdf, Presented at the Theory and Simulation of Disruptions Workshop, 2017.
- [52] R. Samulyak, S. Yuan, N. Naitlho, and P. Parks, "Lagrangian particle model for 3D simulation of pellets and SPI fragments in tokamaks", *Nuclear Fusion*, vol. 61, p. 046 007, 2021. DOI: 10.1088/1741-4326/abdcd2.
- [53] D. Shiraki, J. Herfindal, L. R. Baylor, E. M. Hollmann, C. Lasnier, I. Bykov, N. Eidietis, R. Raman, R. Sweeney, U. Sheikh, S. Gerasimov, S. Jachmich, M. Lehnen, J. Kim, J. J. Jang, S. Meitner, and T. Gebhart, *Particle assimilation during shattered pellet injection*, https://conferences.iaea.org/event/ 217/contributions/16713/, Presented at the IAEA Technical Meeting on Plasma Disruption and their Mitigation, 2020.
- [54] P. Parks, "Modeling dynamic fracture of cryogenic pellets", General Atomics, Tech. Rep., 2016. DOI: 10.2172/1344852.
- [55] N. Mott and E. Linfoot, "A theory of fragmentation", in *Fragmentation of Rings and Shells: The Legacy of N.F. Mott.* Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, p. 207, ISBN: 978-3-540-27145-1. DOI: 10.1007/978-3-540-27145-1_9.
- [56] D. Hu, E. Nardon, M. Lehnen, G. Huijsmans, and D. van Vugt and, "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. DOI: 10.1088/1741-4326/aae614.
- [57] A. Matsuyama, E. Nardon, M. Honda, and M. Lehnen, *ITER SPI modelling for runaway electron avoidance*, https://conferences.iaea.org/event/217/contributions/16702/, Presented at the Theory and Simulation of Disruptions Workshop, 2017.

- [58] E. Nardon, A. Matsuyama, and M. Lehnen, On the possible injection schemes with the ITER SPI system, https://conferences.iaea.org/event/217/ contributions/16702/, Presented at the IAEA Technical Meeting on Plasma Disruption and their Mitigation, 2020.
- [59] P. Stangeby, The Plasma Boundary of Magnetic Fusion Devices, ser. Series in Plasma Physics and Fluid Dynamics. Taylor & Francis, 2000, ISBN: 9780750305594.
- [60] H. P. Summers, The ADAS user manual, version 2.6, http://www.adas.ac.uk, 2004.
- B. J. Braams and C. F. F. Karney, "Conductivity of a relativistic plasma", *Physics of Fluids B: Plasma Physics*, vol. 1, p. 1355, 1989. DOI: 10.1063/1. 858966.
- [62] A. B. Rechester and M. N. Rosenbluth, "Electron heat transport in a tokamak with destroyed magnetic surfaces", *Phys. Rev. Lett.*, vol. 40, p. 38, 1 1978. DOI: 10.1103/PhysRevLett.40.38.
- [63] L. Hesslow, O. Embréus, M. Hoppe, T. 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, p. 905 840 605, 2018. DOI: 10.1017/S0022377818001113.
- [64] S. T. Beliaev and G. I. Budker, *Soviet Physics-Doklady*, vol. 1, p. 218, 1956.
- [65] H. W. Koch and J. W. Motz, "Bremsstrahlung cross-section formulas and related data", *Rev. Mod. Phys.*, vol. 31, p. 920, 4 1959. DOI: 10.1103/RevModPhys. 31.920.
- [66] 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. 074 010, 2018. DOI: 10.1088/1361-6587/aac33e.
- [67] C. F. Karney, "Fokker-Planck and quasilinear codes", Computer Physics Reports, vol. 4, p. 183, 1986, ISSN: 0167-7977. DOI: 10.1016/0167-7977(86) 90029-8.
- [68] K. Särkimäki, O. Embreus, E. Nardon, T. Fülöp, and J. contributors, "Assessing energy dependence of the transport of relativistic electrons in perturbed magnetic fields with orbit-following simulations", *Nuclear Fusion*, vol. 60, p. 126050, 2020. DOI: 10.1088/1741-4326/abb9e9.
- [69] A. Pshenov, A. Kukushkin, E. Marenkov, and S. Krasheninnikov, "On the role of hydrogen radiation absorption in divertor plasma detachment", *Nuclear Fusion*, vol. 59, p. 106 025, 2019. DOI: 10.1088/1741-4326/ab3144.
- [70] D. Morozov, E. Baronova, and I. Senichenkov, "Impurity radiation from a tokamak plasma", *Plasma Physics Reports*, vol. 33, p. 906, 2007. DOI: 10.1134/ S1063780X07110037.
- [71] V. Lukash, A. Mineev, and D. Morozov, "Influence of plasma opacity on current decay after disruptions in tokamaks", *Nuclear Fusion*, vol. 47, p. 1476, 2007. DOI: 10.1088/0029-5515/47/11/009.

- [72] 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, p. 905870207, 2021. DOI: 10.1017/S0022377820001592.
- [73] E. Nardon, J. Artola, C. Sommariva, M. Hoelzl, D. Hu, and G. Huijsmans, JOREK simulations of MGI-triggered disruptions in JET, https://ft.nephy. chalmers.se/?p=abstract&id=11, Presented at the 8th Runaway Electron Modelling (REM) meeting, 2020.
- [74] L. Johnson and E. Hinnov, "Ionization, recombination, and population of excited levels in hydrogen plasmas", *Journal of Quantitative Spectroscopy* and Radiative Transfer, vol. 13, p. 333, 1973, ISSN: 0022-4073. DOI: https: //doi.org/10.1016/0022-4073(73)90064-2.
- [75] D. Verner, E. Verner, and G. Ferland, "Atomic data for permitted resonance lines of atoms and ions from H to Si, and S, Ar, Ca, and Fe", Atomic Data and Nuclear Data Tables, vol. 64, p. 1, 1996, ISSN: 0092-640X. DOI: https: //doi.org/10.1006/adnd.1996.0018.
- [76] M. J. Seaton, "Radiative Recombination of Hydrogenic Ions", Monthly Notices of the Royal Astronomical Society, vol. 119, p. 81, 1959, ISSN: 0035-8711. DOI: 10.1093/mnras/119.2.81. eprint: https://academic.oup.com/mnras/article-pdf/119/2/81/8078011/mnras119-0081.pdf.
- [77] K. Gál, T. Fehér, H. Smith, T. Fülöp, and P. Helander, "Runaway electron generation during plasma shutdown by killer pellet injection", *Plasma Physics* and Controlled Fusion, vol. 50, p. 055 006, 2008. DOI: 10.1088/0741-3335/ 50/5/055006.

Discretisation of SPI deposition kernels

The discretised delta source function, averaged over time step and grid cell, can be expressed as

$$H(r_i, \rho_{\rm p}) = \frac{\min\left(r_{\rm f, i+1}, \rho_{\rm p, max} - \max\left(r_{\rm f, i}, \rho_{\rm p, min}\right)\right)}{V'(\rho_{\rm p, max} - \rho_{\rm p, min})\Delta r_i}.$$

$$\Theta(\rho_{\rm f, i+1} - \rho_{\rm p, min})\Theta(\rho_{\rm p, max} - \rho_{\rm f, i}),$$
(A.1)

where r_i and Δr_i are the discretised radial grid points and grid step size, $\rho_{p,max} = \max(\rho_p(t) \text{ and } \rho_p(t + \Delta t)), \rho_{p,min} = \min(\rho_p(t), \rho_p(t + \Delta t)), \Theta$ is the Heaviside step function. The volume between two flux surfaces separated by an infinitesimal distance dr is given by v'dr. In a cylindrical geometry, we have $V' = \partial V/\partial r = A_{flt} = 4\pi^2 R r$ with cylindrical geometry. The subscript f on r denotes the radial coordinate of the cell faces, $r_{f,i} = r_i - \Delta r_i/2, r_{f,i+1} = r_i + \Delta r_i/2$.

Equation (A.1) may be understood as a box function over the radii passed during the current time step, averaged over the discretised flux tubes. This expression is however not applicable when the shard passes the point of closest approach to the magnetic axis, where the radial coordinate varies non-monotonically along the shard trajectory. The time step where the shard passes closest to the magnetic axis is therefore divided into two parts, one part before and one part after passing the closest approach to the center, and equation (A.1) is then used for these parts separately. To see if the point of closest approach has been passed during the current time step, we check whether the projection of the shard trajectory on the gradient of the radial coordinate has changed sign, i.e. if

$$[\mathbf{x}_{\mathrm{p}}(t) \cdot \nabla r(\mathbf{x}_{\mathrm{p}}(t))] \cdot [\mathbf{x}_{\mathrm{p}}(t + \Delta t) \cdot \nabla r(\mathbf{x}_{\mathrm{p}}(t + \Delta t))] < 0.$$
(A.2)

The value for ρ_p at closest approach can then be found on the line between $\mathbf{x}_p(t)$ and $\mathbf{x}_p(t + \Delta t)$ according to

$$\rho_{\rm p,closest} = |\mathbf{x}_{\rm p}(t) + \delta[\mathbf{x}_{\rm p}(t + \Delta t) - \mathbf{x}_{\rm p}(t)]|, \qquad (A.3)$$

$$\delta = \frac{\mathbf{x}_{\mathrm{p}}(t) \cdot [\mathbf{x}_{\mathrm{p}}(t + \Delta t) - \mathbf{x}_{\mathrm{p}}(t)]}{|\mathbf{x}_{\mathrm{p}}(t + \Delta t) - \mathbf{x}_{\mathrm{p}}(t)|^{2}}.$$
(A.4)

The discretised Gaussian deposition kernel, integrated over the grid cells, can be

expressed in terms of the error function, erf, according to

$$H(r_{i}, \rho_{\rm p}) = \frac{1}{2\pi^{2}R(r_{\rm f,i+1}^{2} - r_{\rm f,i}^{2})} \left(\frac{\operatorname{erf}[(r_{\rm f,i+1} - \rho_{\rm p})/r_{\rm cld}] - \operatorname{erf}[(r_{\rm f,i} - \rho_{\rm p})/r_{\rm cld}]}{2} + \frac{\operatorname{erf}[(-r_{\rm f,i+1} - \rho_{\rm p})/r_{\rm cld}] - \operatorname{erf}[(-r_{\rm f,i} - \rho_{\rm p})/r_{\rm cld}]}{2} \right)$$
(A.5)

The second term accounts for that the deposition profile crosses each flux tube twice, once on each side of the magnetic axis. Note, that this expression is not averaged over the time step, and is therefore only valid if the distance traveled during one time step is small compared to the width $r_{\rm cld}$ of the shielding cloud.
В

Opacity for deuterium radiation

Here we estimate how a large fraction of the line radiation from deuterium, due to excitation and recombination, is expected to be trapped due to opacity of the plasma in the vicinity of the Lyman lines, in a plane partially ionized plasma slab. We follow the model described in ref. [70], which has also been used previously in ref. [71] to study the effect of opacity to impurity radiation during mitigated disruptions. The fraction of trapped radiation is determined mainly by two quantities: the optical thickness of the plasma to the line radiation, and the rate of collisional quenching of excited states. If the rate of collisional quenching is low, even an optically thick plasma would be effectively transparent, as the absorbed photons would then likely be re-emitted again rapidly, and so soon escape from the plasma.

The optical thickness of a plasma slab of thickness h is given by $\tau = k_z^j h$, with k_z^j being the inverse mean free path of the photon emitted by deexcitation from the j^{th} excited state to the ground state of a nucleus of charge z (transitions between excited states are neglected). We will limit ourselves to the first nine excited states due to availability of the necessary data. To further motivate this truncation, we verified that including the ninth excited state only altered the results by the order of 1% compared to only including the first eight excited states. If the natural broadening of the line profile is γ and the net external broadening is Γ , the photon mean free path is given by

$$k_{z}^{j} = \frac{n_{z}(\lambda_{z}^{j})^{2}}{4\pi} \frac{1}{1 + \Gamma/\gamma}.$$
 (B.1)

Here, λ_z^j is the wavelength of a photon emitted by deexcitation from the j^{th} excited state to the ground state of a nucleus of charge z. For a Doppler broadened line, we have

$$\Gamma/\gamma = 1.11 \cdot 10^{10} E_z^{0j} \sqrt{\frac{m_p}{m_i} T_i} \nu_z^j,$$
(B.2)

where m_p is the proton mass, m_i is the ion (or atom) mass, T_i is the ion temperature, E_z^{0j} is the energy difference between the ground state and the j^{th} excited state of an ion with charge z and ν_z^j is the corresponding natural radiative decay rate. In this section, all temperatures and energies are here given in eV, and transition rates are given in s⁻¹.

The effect of collisional quenching of excited states is determined by the ratio of the collisional quenching probability and the radiative decay probability,

$$\beta_z^j = \frac{2.7 \cdot 10^{-13} n_e}{(E_z^{0j})^3 \sqrt{T_e}} \left[1 - \frac{E_z^{0j}}{T_e} \exp\left(\frac{E_z^{0j}}{T_e}\right) E_1\left(\frac{E_z^{0j}}{T_e}\right) \right],\tag{B.3}$$

where $E_1(x) = \int_x^\infty \exp(-t)/t dt$ is the integral exponent. The fraction of radiation escaping from a plane ion slab for a line emitted by deexcitation of the j^{th} excited state from a nucleus of charge z is then given by

$$B_z^j = \frac{1}{\beta_z^j + 1} W_z^j \tag{B.4}$$

where

$$W_{z}^{j} = \frac{(1+\beta_{z}^{j})P_{a}}{\beta_{z}^{j} + P_{a}}$$
(B.5)

and

$$P_{\rm a} = \frac{1}{1 + \tau \sqrt{\pi \ln{(\tau + 1)}}} \tag{B.6}$$

is the probability for a photon to travel a distance h without being absorbed (the "1" in the denominator is included *ad-hoc* in the expression valid in the optically thick limit, to make sure that $P_{\rm a} \rightarrow 1$ as it should for thin plasma slabs).

To find the fraction of the total radiation that escapes the plasma, we also need to know the relative intensity of the different lines. The intensity distribution over the different lines is very different for excited states populated by excitations from lower states (dominantly the ground state) compared to excited states populated by recombination. The reason for this is that excitations from lower states primarily populate the lower excited states, while recombination populates higher excited states to a substantially larger extents. In addition, part of the potential energy change is radiated away during a free-bound transition, resulting in a continuous spectrum to which the plasma is essentially transparent.

For radiation following excitation from lower states, we calculate the relative line intensities based on data found in ref. [74], according to

$$L_z^j \propto E_z^{0j} \phi_z^j, \tag{B.7}$$

where

$$\phi_z^j = r_z^j \bar{n}_{\text{Saha}}(j) \nu_z^j = r_z^j (j+1)^2 \exp\left(E_z^{j\infty}/T_e\right) \nu_z^j \tag{B.8}$$

is proportional to the relative occupation \bar{n}_{Saha} of the j^{th} excited state at Saha equilibrium, and $E_z^{j\infty}$ is the ionization energy from the j^{th} excited state. The coefficients r_z^j describe the deviation from the Saha equilibrium prevailing in a dense enough plasma to reach local thermodynamic equilibrium. These coefficients are tabulated in ref. [74], and the transition rates ν_z^j are tabulated in ref. [75]. The fraction of line radiation following excitations from the ground state escaping the plasma can now be expressed as

$$f_{\rm esc} = \frac{\sum_{z} n_{z} \sum_{j} B_{z}^{j} L_{z}^{j}}{\sum_{z} n_{z} \sum_{j} L_{z}^{j}} = \frac{\sum_{j} B_{0}^{j} E_{0}^{0j} \phi_{0}^{j}}{\sum_{j} E_{0}^{0j} \phi_{0}^{j}}, \tag{B.9}$$

where the second equality holds for hydrogen isotopes where there is only one radiating charge state.

When the excited states are populated by recombination, energy will also be radiated away during the free-bound transition populating the excited bound states, in addition to the subsequent deexcitation through bound-bound transitions. We consider two types of recombination events, one with transitions directly from a free state to the bound ground state, and one consisting of a free-bound transition to an excited bound state, followed by a direct transition to the ground state. If the recombination is radiative, a transition from state j (which can be the ground state, j = 0) will be preceded by a transition from the continuum to state j, while releasing an average photon of energy $E_{z,rec}^j \approx E_z^{\infty 0} + T_e - E_z^{0j}$. The free-bound transitions give rise to a continuus spectrum, to which the plasma is be essentially transparent. The resulting escape factor for recombination radiation becomes

$$f_{\rm esc} = \frac{\sum_{z} n_{z} \sum_{j} (B_{z}^{j} L_{z}^{j} + L_{z,rec}^{j})}{\sum_{z} n_{z} \sum_{j} (L_{z}^{j} + L_{z,rec}^{j})} = \frac{\sum_{j} (B_{0}^{j} E_{0}^{0j} + E_{0,rec}^{j}) \phi_{0}^{j}}{\sum_{j} (E_{0}^{0\,\text{inf}} + T_{e}) \phi_{0}^{j}}.$$
 (B.10)

As the coefficients in ref. [74] only apply to recombination events involving an excited bound state, we replace ϕ_0^j in equation (B.10) with the transition rates from ref. [76], which also include radiative recombination directly to the ground state.

The escape factor as a function of the slab thickness h is shown in figure B.1, for h up to the minor radius a = 2 m of an ITER-like plasma. Results are shown for both radiation following excitations from the ground state, calculated using equation (B.9), and for recombination radiation, calculated using equation (B.10). The plasma parameters considered are an electron density $n_e = 10^{20}$ m⁻³, neutral deuterium density $n_D = 4 \cdot 10^{21}$ m⁻³, and temperature $T_e = 1.38$ eV. These plasma parameters are chosen to match the ones for which there are tabulated data in ref. [74], and to be representative for case 3 described in section 5.3). One can see that the line radiation following excitations from the ground state have an escape factor of not more than a few percent, unless the slab thickness is very small. For the recombination radiation, on the other hand, the escaping fraction is of the order of tens of percent even for rather thick plasma slabs.

These results have several limitations that makes them rather uncertain. They neglected transitions between excited states (to which the plasma would be essentially transparent), and other broadening mechanisms than doppler broadening, both of which would increase the escape factor. On the other hand, including three-body recombination would decrease the escape factor compared to the purely radiative recombinations considered here. If three-body recombination is significant, part of $E_{z,rec}^{j}$ will be absorbed by a second free electron in some of the transitions, which lowers the escape factor. Expressions for the various contributions to the recombination rate found in ref. [70] suggest that three-body recombination and radiative recombination can be comparable at the conditions of interest. Effects of geometry and non-homogeneous plasma profiles are also neglected, as are effects of charge exchange that could also be important according to [70]. For sufficiently large injections, the increase in the free electron density could also be enough to make opacity effects important for neon impurities, depending on how substantial recombination is obtained. Despite these limitations, however, one can conclude that the escape factor is likely to be not more than a few percent for line radiation following excitations from the ground state, but tens of percent for recombination radiation.



Figure B.1: Escaping fraction of deuterium radiation as a function of slab thickness, for a plasma with electron density $n_e = 10^{20} \text{ m}^{-3}$, neutral deuterium density $n_D = 4 \cdot 10^{21} \text{ m}^{-3}$, and temperature $T_e = 1.38 \text{ eV}$. The curves show the escaping fraction for radiation following excitation from the ground state (black solid), calculated by equation (B.9), and for recombination radiation (dotted red), calculated with equation (B.10).

C

Impact of the deposition process features

Once the pellet material is ablated, it starts a homogenisation and equilibration process over the flux surfaces. This process is governed by the pressure gradient along the field lines, the magnetic shear and the shear caused by the pellet-induced poloidal rotation, as discussed in section 3.5.2. The time scale for this process is about 1 ms, which is comparable to the time between the arrival of the first and last pellet shards at a given flux surface. Therefore, the finite time scale of the expansion process might have a significant effect on the SPI dynamics. Moreover, during this process, the pellet material is also subject to an $E \times B$ -drift causing it to drift towards the low field side. Therefore, in reality, the expansion process is not localised to a single flux surface.

Apart from the direct displacement on the final deposition profile due to the cross-field drift, the final deposition profile is also indirectly affected by the finite expansion time. The reason for this is that the response of the plasma to the pellet material affects any further ablation. For pellets with a significant impact on the plasma temperature, the interaction between the plasma and the pellet material is self-regulating; when the pellet material is ablated, it instantaneously cools the plasma by dilution, resulting in a slowing down of the ablation. This self-regulation is enhanced by assuming an instantaneous and local homogenisation and equilibration process. However, the finite expansion time scale and the cross-field drift reduces the self-regulation by delaying and displacing the cooling. This might have a significant effect on the ablation, and hence also the final deposition profile.

In this section, we investigate the sensitivity to a finite expansion time scale and cross-field drift on the final deposition profile. To do this, we emulate the effect that a shard does not experience instant dilutive cooling of the background plasma by the vanguard shards, by depositing the pellet material one radial grid cell behind the ablating shards. Such an artificial shift can be used to roughly mimic the effect of both radial drifts and expansion time delays, as the shift also introduces a delay between the passing of the shards and the deposition of material on a given flux surface. The duration of this time delay would be the width of the shift divided by $\langle v_{\rm p} \rangle$. When a time delay due to the finite expansion time is involved, the thermal energy going into the pellet cloud is not instantaneously returned to the background plasma. In such a case, the shards affect the local background plasma temperature by absorbing the heat flux into the pellet cloud. When emulating the effect of a

finite expansion time, we account for this energy absorption by adding a term

$$\left(\frac{\partial W_{\rm M}}{\partial t}\right)_{\rm abs} = -\sum_{k=1}^{N_{\rm s}} 2q_{\rm tot} \pi r_{\rm cld}^2 H(r,\rho_{\rm p,k}),\tag{C.1}$$

to the energy transport equation (4.8). For simplicity, when including equation (C.1), the delayed returning of the energy absorbed in the pellet cloud is neglected. Also note that the finite $r_{\rm cld}$ is only used to determine the magnitude of the heat absorption, while the deposition kernel $H(r, \rho_{\rm p,k})$ is still assumed to be a delta function, unless otherwise noted. This is motivated by that a realistic $r_{\rm cld}$ of the order of 1 cm (see section 3.5.1) is smaller than the radial resolution used in our simulations.

In addition to the shift of the deposition kernel discussed above, we also estimate the effect of varying the shape of the deposition kernel. For the simulations presented in chapter 5, the deposition kernel was assumed to be a delta function in the continuous space. When discretised, however, the lower limit for the width of the kernel is set by the width of the radial grid cell. The effect of this finite kernel width is investigated here by varying the number $N_{\rm r}$ of radial grid points. To further study the effect of different deposition kernel shapes, we also investigate the effect of assuming a Gaussian deposition kernel, $H(r, \rho_{\rm p,k}) \propto \exp(-(r - \rho_{\rm p,k})^2/r_{\rm cld}^2)$. Such a kernel has been used for SPI-studies in ITER-like plasmas previously in e.g. [24]. Although a realistic value for $r_{\rm cld}$ is of the order of 1 cm, for numerical reasons, the value used for $r_{\rm cld}$ in ref. [24] is 20 cm. Therefore, we also consider such a large value of $r_{\rm cld}$ and compare the final deposition profile to the one obtained with a delta-like kernel.

We start with performing some of the studies described above for a single unshattered pellet injection into a JET-like plasma. The assumed tokamak settings are a = 1 m and $R_0 = 2.96$ m, with initial profiles of temperature $T_{\rm M}(t = 0) =$ $T_0(1-0.75(r/a)^2)^2$ and density $n_{\rm M}(t=0) = n_0(1-0.9(r/a)^2)^{2/3}$, where $T_0 = 3.1$ keV and $n_0 = 2.8 \cdot 10^{19} \text{ m}^{-3}$. The injected pellet consists of pure deuterium and has a radius of 3 mm, containing $6.8 \cdot 10^{21}$ particles. The speed of the pellet is $v_{\rm p} = 160$ m/s, directly aimed towards the plasma core ($\alpha_{\rm p} = 0$). When applying a shift to the deposition kernel as discussed above, the heat absorption in the pellet cloud is calculated assuming $r_{\rm cld} = 1$ cm. These injection and plasma settings are the same as those used in section 4.1 in ref. [77]. The simulations in ref. [77] are similar to our shifted model, in the sense that the background density is assumed to be unaffected by the pellet material as far as the pellet ablation is concerned. Thus, the pellet ablation is not affected by dilutive cooling of the background plasma. Note, however, that the pellet model used in ref. [77], which is described in detail in ref. [46], accounts for the additional shielding of the pellet by the ablated but not yet homogenised and equilibrated material, which is not accounted for by our shift-approximation.

Figure C.1 a) shows simulations of this case, demonstrating the impact of shifting the deposition profile. The result using two different radial resolutions are compared. For the cases with an instantaneous and local deposition (labeled "local"), the pellet passes through the center of the plasma. This is not observed in ref. [77] (see their figure 1 c). When applying the shift, we see the pellet burns out completely about half way through the plasma, which is qualitatively similar to figure 1 c) in ref. [77]. The penetration depth is however somewhat smaller here, which is expected because the pellet code used in ref. [77] also accounts for the additional shielding of the pellet by the ablated but not yet homogenised and equilibrated material. Similar results are obtained when changing the number of radial grid cells $N_{\rm r}$ from 11 to 50, confirming that the radial spread only has a minor impact on the final deposition profile.



Figure C.1: Deposition profiles obtained assuming an instantaneous and local flux surface homogenisation and equilibration (labeled "local"), and applying a shift of the deposition kernel one grid cell behind the pellet, to mimic the effect of a finite expansion time and the cross field drift (labeled "shifted"). The radial grid spacing and radial distribution of the ablated material is also varied. a) A JET-like case with an unshattered deuterium pellet injection with $N_{\rm inj} = 6.8 \cdot 10^{21}$ and $v_{\rm p} = 160$ m/s, similar to the case studied in section 4.1 in ref. [77]. b) An ITER-like deuterium SPI case similar to the one shown in ref. [58] (which uses the same initial profiles as used in section 5.1-5.3), with $N_{\rm inj} = 2.2 \cdot 10^{24}$, $N_{\rm s} = 300$ and $\langle v_{\rm p} \rangle = 200$ m/s.

We now move on to a similar comparison between the effect of a local and shifted deposition model for an SPI case in an ITER-like plasma. We also consider the impact of using an unrealistically wide Gaussian deposition kernel, similarly to what is used in ref. [24]. Here, the tokamak settings and initial profiles are the same as the ones used in this work in section 5.1-5.3. The pellet is made of pure deuterium, and the injection parameters used here are $N_{\rm s} = 300$ and $\langle v_{\rm p} \rangle = 200$ m/s (i.e. a somewhat slower pellet speed than in chapter 5). These settings are the same as the ones used for the pure deuterium case studied with the code INDEX in ref. [58]. The model for the pellet ablation and the resulting effect on the background plasma in this code is similar to the model used in our work, presented in section 4.1. The result is shown in figure C.1 b). When using local deposition, rather good agreement with [58] is seen, regardless of the shape of the radial spread of the ablated material (note, here we show the final profile, while the latest time slice shown in ref. [58] is at 6 ms. At

this time, about 2% of the pellet material is still not ablated, which introduces a small discrepancy for the core penetration). Note that the heat absorption is only included in the shifted cases; for the green dashed curve, $r_{\rm cld}$ is only used as the width of the Gaussian deposition kernel.

When applying a shift to the deposition, a shift of one radial grid cell with $N_{\rm r} = 11$ in a plasma with a = 2 m would approximately correspond to a time delay of $a/(N_{\rm r}\langle v_{\rm p}\rangle) \approx 1$ ms, using $\langle v_{\rm p}\rangle = 200$ m/s. Such a time scale is similar to that expected for the flux surface homogenisation of the pellet material [43]. As the number of shards increase, the total heat absorption increases. This happens because the radius of the pellet cloud absorbing the heat flux is mostly determined by the distance the particles travel before they ionize, which is not strongly affected by the shard size. Depending on the pellet cloud radius, this absorption may compensate substantially for the lack of dilutive cooling in the shifted cases.

Moreover, using a variance of $\Delta v_{\rm p} = 0.2 \langle v_{\rm p} \rangle$, the radial extension of the plume of shards becomes wider than the applied shift of $a/N_{\rm r} = 20$ cm after about 2.5 ms. During this time, the shards travel on average about a fourth of the distance towards the plasma core. Thus, the later arriving shards still have time to feel the dilutive cooling by the earlier arriving shards. As a result, depending on the value of $r_{\rm cld}$, the final profile might not be very different from those obtained assuming a local deposition and instantaneous expansion over the flux surfaces. This can be seen by comparing the red dotted curve with the blue dash-dotted curve in figure C.1 b), where the latter assumes $r_{\rm cld} = 2$ cm when calculating the heat absorption. The discrepancy could be further reduced by accounting for the additional shielding by the ablated but not yet homogenised and equilibrated material. However, the solid black curve in figure C.1 b) with $r_{\rm cld} = 1$ cm differs substantially from the local cases. In conclusion, a more accurate and self-consistent model for the transient plasma response to the deposited pellet material would be needed to make quantitatively accurate predictions for the final deposition profile.